

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



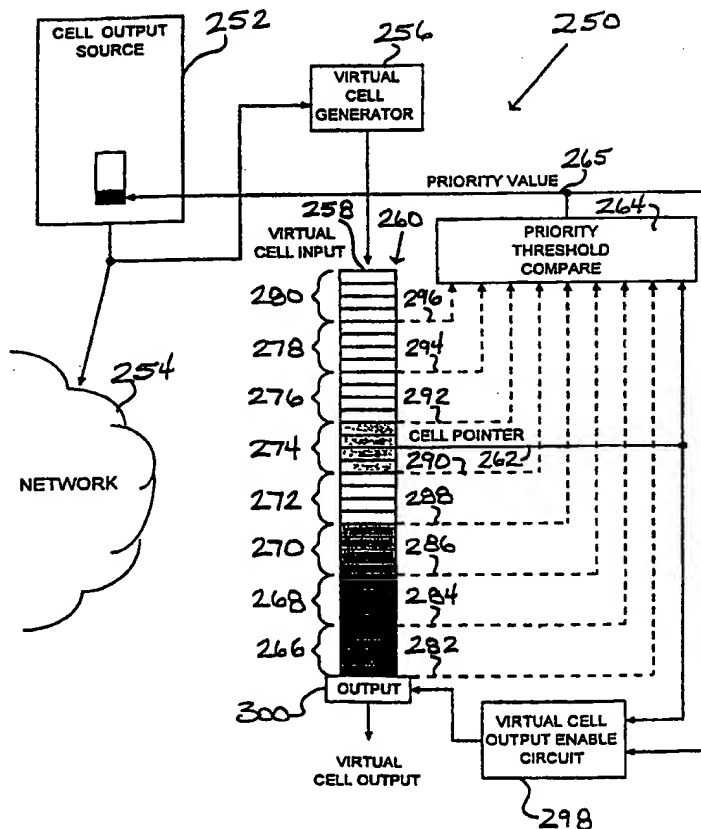
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04Q 11/04, H04L 12/56		A1	(11) International Publication Number: WO 00/57668
			(43) International Publication Date: 28 September 2000 (28.09.00)
(21) International Application Number: PCT/FI99/00232		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 24 March 1999 (24.03.99)			
(71) Applicant (for all designated States except US): NOKIA NETWORKS OY [FI/FI]; Keilalahdentie 4, FIN-02150 Espoo (FI).			
(72) Inventors; and			
(75) Inventors/Applicants (for US only): RUUTU, Jussi [FI/FI]; Ilansuu 2 D 4, FIN-02210 ESPOO (FI). SARANKA, Sari [FI/FI]; Tonttutyönpoiku 4 A, FIN-02200 Espoo (FI). KILKKI, Matti [FI/US]; Apartment 45, 425 Woburn Street, Lexington, MA 02420 (US).			
(74) Agent: JOHANSSON, Folke; Nokia Corporation, P.O. Box 319, FIN-00045 Nokia Group (FI).		Published With international search report.	

(54) Title: BUFFER-BASED TRAFFIC MEASUREMENT SYSTEM AND METHOD FOR NOMINAL BIT RATE (NBR) SERVICE

(57) Abstract

A system and method for controlling the cell priority of cells transmitted from a cell source to network destinations across a network connection is provided. Each of the cells transmitted is registered in a measurement buffer upon being output from the cell source. The registered cells in the measurement buffer are discharged at a discharge rate corresponding to the current occupancy level of the measurement buffer. Cell priority values are generated by correlating the current buffer occupancy level to one of a plurality of predetermined buffer occupancy ranges to determine which occupancy range coincides with the current occupancy level. The generated cell priority values are loaded into cells transmitted from the cell source for use by the network connection.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

**BUFFER-BASED TRAFFIC MEASUREMENT SYSTEM AND METHOD FOR
NOMINAL BIT RATE (NBR) SERVICE**

5

FIELD OF THE INVENTION

The present invention relates generally to
10 communications systems, and more particularly, to a
buffer-based system and method for measuring cell
transmission rates and assigning cell priorities to
cells sent through a network employing a nominal bit
rate (NBR) service.

15

BACKGROUND OF THE INVENTION

The state of communications technology,
20 particularly that which affects the Internet, is
currently in flux and subject to rapid and often
uncoordinated growth. The ubiquity and variety of
personal computers and set-top boxes has placed
significant pressure on the providers of communications
25 system infrastructure to accommodate the alarming
increase in the number of new users that demand
immediate access to Internet and other network
resources. The rapid development of new and
sophisticated software made available to users of such
30 services places additional demands on system
infrastructure.

Conducting commerce over the Internet and
other networks is a practice that is gaining acceptance
and popularity. By way of example, traditional on-line
35 services, such as those offered by Internet providers,
typical charge customers a monthly fee for access to

basic services and resources, such as proprietary and public databases of information. Such traditional service providers also advertise any number of products or services which are purchasable on-line by the user.

5 Other forms of Internet commercialization currently being considered or implemented include offering of video and audio conferencing services, and a variety of other real-time and non-real-time services. The providers of these services, as well as the
10 providers of communications system infrastructure, are currently facing a number of complex issues, including management of network capacity, load, and traffic to support real-time, non-real-time, and high-bandwidth services, and implementing a viable billing scheme that
15 accounts for the use of such services.

 The communications industry is expending considerable attention and investment on one particular technology, referred to as asynchronous transfer mode (ATM), as a possible solution to current and anticipated
20 infrastructure limitations. Those skilled in the art understand ATM to constitute a communications networking concept that, in theory, addresses many of the aforementioned concerns, such as by providing a capability to manage increases in network load,
25 supporting both real-time and non-real-time applications, and offering, in certain circumstances, a guaranteed level of service quality.

 A conventional ATM service architecture typically provides a number of predefined quality of
30 service classes, often referred to as service categories. Each of the service categories includes a number of quality of service (QoS) parameters which define the nature of the respective service category. In other words, a specified service category provides
35 performance to an ATM virtual connection (VCC or VPC) in

a manner specified by a subset of the ATM performance parameters. The service categories defined in the ATM Forum specification referenced hereinbelow include, for example, a constant bit rate (CBR) category, a real-time
5 variable bit rate (rt-VBR) category, a non-real-time variable bit rate (nrt-VBR) category, an unspecified bit rate (UBR) category, and an available bit rate (ABR) category.

The constant bit rate service class is
10 intended to support real-time applications that require a fixed quantity of bandwidth during the existence of the connection. A particular quality of service is negotiated to provide the CBR service, where the QoS parameters include characterization of the peak cell
15 rate (PCR), the cell loss rate (CLR), the cell transfer delay (CTD), and the cell delay variation (CDV). Conventional ATM traffic management schemes guarantee that the user-contracted QoS is maintained in order to support, for example, real-time applications, such as
20 circuit emulation and voice/video applications, which require tightly constrained delay variations.

The non-real-time VBR service class is intended to support non-real-time applications, where the resulting network traffic can be characterized as
25 having frequent data bursts. Similarly, the real-time variable bit rate service category may be used to support "bursty" network traffic conditions. The rt-VBR service category differs from the nrt-VBR service category in that the former is intended to support real-
30 time applications, such as voice and video applications. Both the real-time and non-real-time VBR service categories are characterized in terms of a peak cell rate (PCR), a sustainable cell rate (SCR), and a maximum burst size (MBS).

The unspecified bit rate (UBR) service category is often regarded as a "best effort service," in that it does not specify traffic-related service guarantees. As such, the UBR service category is
5 intended to support non-real-time applications, including traditional computer communications applications such as file transfers and e-mail.

The available bit rate (ABR) service category provides for the allocation of available bandwidth to
10 users by controlling the rate of traffic through use of a feedback mechanism. The feedback mechanism permits cell transmission rates to be varied in an effort to control or avoid traffic congestion, and to more effectively utilize available bandwidth. A resource
15 management (RM) cell precedes the transmission of data cells, which is transmitted from source to destination and back to the source, in order to provide traffic information to the source.

Although the current ATM service architecture
20 described above would appear to provide, at least at a conceptual level, viable solutions to the many problems facing the communications industry, ATM, as currently defined, requires implementation of a complex traffic management scheme in order meet the objectives
25 articulated in the various ATM specifications and recommendations currently being considered. In order to effectively manage traffic flow in a network, conventional ATM traffic management schemes must assess a prodigious number of traffic condition indicators,
30 including service class parameters, traffic parameters, quality of service parameters and the like. A non-exhaustive listing of such parameters and other ATM traffic management considerations is provided in ITU-T Recommendation I.371, entitled Traffic Control and
35 Congestion Control in B-ISDN, and in Traffic Management

Specification, version 4.0 (af-tm-0056.000, April 1996), published by the Technical Committee of the ATM Forum.

One traffic parameter utilized in conventional traffic management schemes includes the rate at which
5 cells are transmitted by the user. Conventional ATM traffic management principles provide for cell transmission rate adjustments based on feedback control information received from the network. This additional complexity may be undesirable, as it may be faster and
10 more efficient to perform these operations at the user/network interface (UNI). Furthermore, these conventional measurement principles provide feedback information periodically rather than continually providing cell rate information at the UNI. The
15 conventional traffic management schemes also fail to provide measurement techniques for multiple cell priority values as is utilized in connection with the nominal bit rate service.

Accordingly, there is a need in the
20 communications industry for a network arrangement and method that provides for a traffic management system and method which is conceptually uncomplicated to implement. A further need exists for a system and methodology that provides a buffer-based cell traffic measurement
25 principle at the UNI, which allows cell rate information to be continuously available to the user rather than only providing information periodically, and which further avoids the complexity of cell feedback techniques. A system and methodology for assigning
30 multiple priority levels, rather than a general on/off priority designation, is also desirable for use in connection with a nominal bit rate service as described herein. The present invention fulfills these and other needs, and offers other advantages over the prior art
35 traffic management approaches.

SUMMARY OF THE INVENTION

The present invention is directed to a system
5 and method for controlling the priority assigned to
information elements transmitted across a network
connection. A buffering scheme is employed to measure
the transmission rate of information elements, such as
cells, and to assign one of a sequential collection of
10 cell priorities to cells transmitted from a cell source.
The cells, each loaded with one of the cell priorities,
are transmitted via a network connection having a
quality of service based on a nominal bit rate (NBR). A
nominal bit rate represents an expected, but not
15 guaranteed, bit rate associated with a particular user
or connection, and utilizes multiple cell priorities in
managing cell traffic. The present invention provides a
buffer-based system and method to support the cell
traffic management in such a system.

20 In accordance with one embodiment of the
invention, a method for controlling the cell priority of
cells transmitted from a cell source to network
destinations across a network connection is provided.
Each of the cells transmitted is registered in a
25 measurement buffer upon being output from the cell
source. These registered cells in the measurement
buffer are discharged at a discharge rate corresponding
to the current occupancy level of the measurement
buffer. Cell priority values are generated by
30 correlating the current buffer occupancy level to one of
a plurality of predetermined buffer occupancy ranges to
determine which occupancy range coincides with the
current occupancy level. The generated cell priority
values are loaded into cells transmitted from the cell
35 source.

In accordance with another embodiment of the
invention, registering the cells in the measurement

buffer includes providing a virtual cell in the measurement buffer for each of the cells transmitted from the cell source. In one embodiment this involves generating a duplicate cell for each of the cells transmitted from the cell source, and in another embodiment includes incrementing a buffer address pointer by one cell amount for each of the cells transmitted from the cell source. Cells are generally discharged at a rate regulated by providing one of a plurality of output enable signals to the output of the measurement buffer, where a different output enable signal is provided depending on which of the predetermined buffer occupancy ranges coincides with the current occupancy level.

The above summary of the present invention is not intended to describe each illustrated embodiment or implementation of the present invention. This is the purpose of the figures and the associated discussion which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

Other aspects and advantages of the present invention will become apparent upon reading the following detailed description and upon reference to the drawings in which:

Fig. 1 is a flow diagram illustrating a general procedure for communicating cells of information between a user/network interface and a network using a nominal bit rate service in accordance with an embodiment of the present invention;

Fig. 2 is a system block diagram of a nominal bit rate service architecture in accordance with an embodiment of the present invention;

Fig. 3 illustrates in greater detail a procedure for transmitting cells of information between a user/network interface and a network using a nominal bit rate service in accordance with another embodiment of the present invention;

Fig. 4 illustrates in flow diagram form a general procedure for filtering cells at a network node in accordance with an embodiment of a nominal bit rate service;

Fig. 5 is a block diagram of an embodiment of a system for filtering cells at a network node in accordance a nominal bit rate service;

Fig. 6 is a block diagram of an alternative embodiment of a system for filtering cells at a network node in accordance with a nominal bit rate service;

Fig. 7 illustrates one embodiment of a network cell in accordance with the NBR service of the present invention;

Fig. 8 is a block diagram illustrating a cell traffic measurement and priority assignment system in accordance with one embodiment of the present invention;

Fig. 9 is a flow diagram illustrating one manner of comparing priority threshold values to the current buffer occupancy levels;

Fig. 10 is a block diagram of one embodiment of a priority threshold compare unit for performing the compare function of the priority threshold compare unit;

Fig. 11 is a graphical representation of the buffer occupancy level versus the rate at which the buffer is emptied where the buffer is implemented as a counter;

Fig. 12 is a graphical illustration of the decreasing buffer occupancy level as a function of time from a full occupancy level where no cell arrivals occur;

Fig. 13 is a graphical representation of the relationship between the buffer occupancy level, the emptying rate, and various buffer discharge times where the buffer occupancy levels are equally spaced;

5 Fig. 14 graphically illustrates the calculation of the number of cells discharged from the buffer between any two cell arrivals;

Fig. 15 is a flow diagram illustrating one embodiment of the invention for the equally spaced
10 buffer occupancy level implementation;

Fig. 16 is a graphical representation of the relationship between the buffer occupancy level, the emptying rate, and various buffer discharge times where the buffer occupancy levels are logarithmically spaced;

15 Fig. 17 is a graphical illustration of the logarithmically decreasing buffer occupancy level as a function of time where no cell arrivals occur;

Fig. 18 illustrates in flow diagram form a general methodology for which the output rate of the
20 virtual cells from the buffer is modified;

Fig. 19 provides one example of a cell output enable circuit in accordance with the present invention; and

25 Figs. 20-21 are graphical depictions of the relationship of average cell loss ratio, P_{loss} , as a function of priority level for four specific load levels.

DETAILED DESCRIPTION OF THE ILLUSTRATED EMBODIMENTS

30

In the following description of the various embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration various embodiments in which the
35 invention may be practiced. It is to be understood that other embodiments may be utilized, and structural and

functional modifications may be made without departing from the scope of the present invention.

The service concept described herein may be viewed as a Simple Integrated Media Access (SIMA) service model. The SIMA service model incorporates the basic properties of ATM with the addition of eight priority levels as defined within the context of a new service concept termed a nominal bit rate service. In general, the NBR service provides for the simple and efficient division of network capacity amongst different connections and the charging of user fees for the use of such connections. A network that embraces a basic version of the SIMA service model does not need to perform many of the traditional and burdensome traffic management functions involving traffic descriptors, quality of service parameters, service classes, connection admission control (CAC), or usage parameter control (UPC).

All of these functions are effectively replaced by functions performed by two autonomous units: a measuring unit, provided at a user/network interface, and a cell scheduling and buffering unit, provided at a network node. The SIMA service concept, from a user's perspective, is simple and understandable, because there are no pre-defined traffic or quality parameters associated with each connection, and charging for connection usage is based solely on the value of NBR and the duration of the connection.

A typical implementation of a SIMA service utilizes two primary components: access nodes and core network nodes, which have fundamentally different functional responsibilities. For example, access nodes, which may be a user/network interface, perform the task of measuring traffic for every connection, whereas at the core network nodes, the traffic control functions do

not need to know anything about the properties of individual connections.

The elegant simplicity of the SIMA service model offers obvious advantages to the manufacturers of infrastructure hardware and software. For example, ATM switches or crossconnects can be built using individual cell scheduling and buffering units, switching fabrics, and routing functions. By using ATM virtual paths or IP switching technology, the routing tasks may be reduced in complexity. In addition, packet discarding and priority feedback features may be included in the cell scheduling and buffering units without negatively impacting their automaticity. Also, simple implementation of network nodes may result in the availability of a relatively inexpensive, high capacity network infrastructure.

The more complex unit of the SIMA service infrastructure concerns the access nodes. Such access nodes will typically include a measuring unit to measure the traffic stream of every connection in real-time, and a computation unit for determining a priority to be assigned to every cell. These additional features should be implementable at a level of difficulty no greater than that for effecting UPC in conventional ATM networks.

The present invention provides a system and methodology for measuring the cell transmission rate from the user/network interface (UNI), and for assigning one of a plurality of priority levels, corresponding to a sequential collection of priority levels, to each of the cells output from the UNI. The system utilizes a buffer-based measuring concept wherein virtual cells corresponding to cells output from the cell output source are input into a buffer. The occupancy level of the buffer provides an indication of the cell

transmission rate of an individual connection. The calculated cell transmission rate is compared to the NBR of the connection to determine how to manipulate cell priority values assigned to the cells.

5 The buffer is implemented such that the cell occupancy level of the buffer increases by one cell at each virtual cell arrival. Cells are discharged at a rate corresponding to the buffer occupancy level. The buffer is divided into multiple buffer occupancy ranges,
10 each of which corresponds to a priority level to be assigned to cells output from the cell output source. Generally, where the buffer occupancy level is high, a high cell loss priority level is assigned to the cells, which indicates that they are more likely to be
15 discarded at a network node. Analogously, a low buffer occupancy level indicates a lower cell transmission rate, and a greater likelihood of successful passage through the nodes on the network connection.

Referring now to Fig. 1, there is shown a
20 general methodology for transmitting information between a user/network interface and a network over an NBR service connection. Initially, a user negotiates or selects 40 a nominal bit rate with the network operator, which may be performed prior to, or at the time of,
25 establishing the connection. In one embodiment, the user informs the network operator that a desired NBR is required, and the requested connection bandwidth is allocated to the user. The network operator, in
accordance with this embodiment, need not perform the
30 task of analyzing current network load conditions existing at the core network nodes prior to establishing or releasing an NBR connection. In an alternative embodiment, the network operator performs the task of determining network load status prior to establishing or
35 releasing an NBR connection, although this task may not

be necessary in a properly dimensioned network supporting an NBR service.

Depending on a particular application, the user selects 42 a real-time or a non-real-time network connection. The process of determining the priority level (PL) of each cell, which indicates the importance or criticality of the cell relative to other cells, involves measuring 44 the actual or measured bit rate (MBR) of the selected real-time or non-real-time connection at the UNI. The priority level of each cell is determined 46 at the UNI. In one embodiment of the invention, a ratio of the MBR and the NBR is used to determine 46 the PL.

After computing the priority level of each cell at the UNI, the cells are transmitted 48 to the network, such as to a node of the network. A network node, upon arrival of a cell transmitted from the UNI, performs a cell filtering process by which the node determines whether to accept or discard a particular cell. The cell filtering process involves determining 50 the state of one or more buffers or memories of the network node to determine a buffer or memory occupancy level. The node accepts or discards 52 a cell based on the priority level of the cell and the state of the node buffer. Cells that meet the filtering criteria determined at the node are accepted, buffered, and eventually transmitted 54 to another node in the network or another network in a manner consistent with the expected quality of service for the connection.

Concerning the embodiment illustrated in block diagram form in Fig. 2, there is shown a user 20 that employs a UNI 24 to communicate with a network 30. The user 20 negotiates a nominal bit rate with the network operator 22. The network operator 22 evaluates the user's NBR request based on a number of factors,

including the NBR negotiated with other users 20, the number and nature of different connections associated with other users of the network, and other factors affecting network capacity and traffic flow. In principle, NBR can be zero, in which case all cells communicated through the UNI 24 are given lowest priority within the network 30. The value of NBR may also be greater than the transmission capacity at the UNI 24. If the value of NBR is significantly greater than the transmission capacity, for example, all cells transmitted from the UNI 24 are given highest priority within the network 30. It is noted that the priority level of a cell as defined herein has meaning within a network or a number of networks that embraces the NBR service concept. Cells that traverse beyond a network that offers an NBR service, such as by use of an network/network interface (NNI), may be processed in accordance with the traffic management strategy employed by such other network.

In contrast to conventional network services which are designed to provide a guaranteed quality of service, the network operator 22 does not guarantee the continuous availability of the user negotiated NBR. A properly dimensioned network, however, should provide adequate bandwidth to virtually ensure, although not guarantee, the availability of an established NBR. It is noted that all users who are transmitting data with an equivalent NBR encounter approximately the same quality of service.

Having established an NBR with the network operator 22, the user 20 is permitted to communicate information to a desired destination 36 via the network 30. A measuring unit 26 measures the actual or instantaneous bit rate (i.e., MBR) of each cell communicated between the UNI 24 and the network 30.

Prior to departure of a cell from the UNI 24, a priority level computing unit 28 determines a priority level for the cell using the negotiated NBR and the MBR. In accordance with one embodiment, one of eight priority levels may be attributed to a given cell. The priority level computing unit 28 determines the priority level of a particular cell by computing a ratio of MBR to NBR. The priority level determined by the computing unit 28 is assigned to the cell which is then transmitted from the UNI 24 to the network 30.

The UNI 24 transmits the cell, which contains priority level information, to a node of the network 30, such as node_A 32. The node_A 32 accepts or discards the cell received from the UNI 24 based on the priority level of the cell and the buffering capacity of node_A 32. In general, as the occupancy of the buffer or memory of node_A 32 increases (i.e., becomes more filled), cells having a lower priority (i.e., higher priority level value) are discarded in favor of accepting cells having a higher priority (i.e., lower priority level value). As the occupancy level of the buffer of node_A 32 decreases (i.e., becomes less filled), the node_A 32 becomes increasingly tolerant toward accepting cells of lower priority (i.e., higher priority level values). Cells that are buffered in node_A 32 are subsequently transmitted to another node in the network 30, such as node_B 34, or other network and, ultimately, to an end-destination 36.

The exemplary network 30 used for purposes of this description is depicted as a network having two intermediary nodes 32 and 34. These nodes represent network data communications elements such as routers, switches and multiplexers. However, as will be appreciated by those skilled in the art, the present invention may likewise be implemented in various multi-

node network structures such as multipoint, star, ring, loop and mesh network topologies used in networks ranging from local area networks (LAN) to proliferative global area networks (GAN) such as the Internet.

5 Figures 3-5 illustrate a procedure for scheduling and buffering cells in accordance with one embodiment of an NBR service methodology. Referring now to Fig. 3, a user establishes 60 an NBR with a network operator. It may be desirable, although not required, 10 to initially set the service class 62 to a non-real-time (nrt) service class as a default setting. Depending on a particular application, the user may require a real-time (rt) service class 64, which may be set by the user directly or, typically, by the user's application or 15 communications software. If the user requires a real-time connection, each cell transmitted from the user's UNI will have the service class bit in the cell header set to indicate that the payload of the cell contains real-time information 70. It is noted that within the 20 context of a network implemented in accordance with the NBR concept of the present invention, real-time service class connections are expected to support virtually any real-time application without the need to specify particular cell transfer delay (CTD) and cell delay 25 variation (CDV) parameters. As such, the conventional procedure of setting CTD and CDV bits of the cell header to appropriate values to accommodate the real-time service requirements of the connection is altogether obviated.

30 If the user does not require a real-time service connection, the default non-real-time service class condition remains operative. As such, the rt/nrt service class bit of each cell header is set to indicate that the payload of the cell includes non-real-time 35 information 66. It is noted that the NBR service

disclosed herein does not utilize the cell loss priority (CLP) scheme used by conventional ATM traffic management approaches. As such, the CLP bit in the cell header may instead be used to discern between real-time and non-real-time payloads.

In the above described embodiment, each cell transmitted over a connection is designated as either a real-time cell or a non-real-time cell, such as by appropriately setting the rt/nrt service class bit of the cell header. In an alternative embodiment, depending on a user's requirements, a connection may be designated as being either a real-time or non-real-time connection, and the cells communicated over such a connection need not be individually assigned a real-time or non-real-time status. Each node for a given connection, for example, may perform a table look up procedure upon arrival of a cell at the node to determine whether the cell is associated with a real-time or a non-real-time connection. Thus, in accordance with this embodiment, a cell header bit need not be reserved for distinguishing between real-time and non-real-time cells.

After the rt/nrt service class header bit has been set in the above-described manner, the actual bit rate of a particular cell to be transmitted between the UNI and the network is measured τ_4 . Since, in practice, the actual bit rate may be subject to significant variability over time, a measuring unit of the UNI employs an averaging measuring principle to determine the actual or instantaneous bit rate, MBR_i .

In general, the UNI measures τ_4 the actual bit rate of a cell, such as cell_i, by approximating the actual or instantaneous bit rate of the connection within a measuring period having a duration that is appropriate for the particular connection (e.g., a real-

time or non-real-time connection). The present invention provides for the measurement of the instantaneous bit rate, MBR_i .

Having determined 74 the measured bit rate, MBR_i , of the i :th cell, the priority level of the i :th cell is computed using the measured bit rate, MBR_i , and the nominal bit rate, NBR . In accordance with one embodiment, it is assumed that a cell may be distinguished from other cells using a cell prioritization scheme that employs eight priority levels. In order to indicate which of the eight priority levels is attributed to a particular cell, each cell allocates three bits for this purpose.

In accordance with current ATM specifications, an ATM cell is specified as a unit of transmission having a fixed-size frame consisting of a 5-octet header and a 48-octet payload. It is appreciated that the necessity to allocate three bits in the cell header for the purpose of designating cell priority level may require utilization of currently defined ATM header bits. By way of example, it may be possible to utilize the current Generic Flow Control (GFC) field which constitutes a total of four bits. In this case, three bits may be allocated for designating cell priority level and one bit may be designated as the rt/nrt service class bit. It may be possible, in accordance with another embodiment, to allocate other header bits for the purpose of indicating one of eight priority levels and rt/nrt service class by deviating from the five-octet header ATM specification.

As such, other header bits may be redefined to represent cell priority level and service class designations. Alternatively, one or more bits required to specify cell priority level and/or service class may be situated outside of the currently defined ATM cell

header. The need to make a minor modification to the existing ATM cell header definition is significantly offset by the substantial advantages offered by employing the NBR service scheme of the present invention, such as a significant reduction in network and traffic management overhead and complexity.

It is understood that the number of priority levels may be less than eight or greater than eight. By way of example, if it is assumed that four cell header bits are allocated for purposes of indicating a cell's priority level, as many as 2^4 (i.e., $2^{n\text{-bits}}$) or 16 priority levels may be defined. Increasing the number of priority levels within the context of an NBR service permits the network operator to make finer adjustments to the bandwidth of a particular connection when managing network traffic. The price for this finer level of traffic control is the additional cell header bit or bits needed to resolve a greater number of priority levels.

A priority level computing unit determines the priority level of each cell, such as cell_i. In accordance with one embodiment of the present invention, and assuming that the measured bit rate is MBR_i when the i:th cell is transmitted to the network, the priority level (PL_i) of cell_i may be calculated using the following equation:

$$x = 4.5 + \frac{\ln(MBR_i / NBR)}{\ln(2)}$$

$$PL_i = \begin{cases} 7 & \text{if } x \geq 7 \\ \lfloor x \rfloor & \text{if } 0 < x < 7 \\ 0 & \text{if } x \leq 0 \end{cases} \quad [1]$$

where, $\lfloor x \rfloor$ represents the integer part of x . The present

invention provides for the determination of cell priority levels (PL_i) using a buffer-based cell measurement and priority level assignment technique.

As will be discussed hereinbelow in accordance with an embodiment in which both NBR and traditional ATM service connections are accommodated, the zero priority level, $PL = 0$, is reserved for those connections that use an ordinary ATM service with guaranteed bandwidth and quality of service. Accordingly, Equation [1] above may be modified so as to produce cell priority levels ranging between $PL = 1$ and $PL = 7$, such that $\lfloor x \rfloor$ represents the integer part of x , if $1 < x < 7$.

It can be seen by application of Equation [1] above that if a connection is exploiting network capacity in excess to the connection's negotiated NBR value, the priority level of cell_i is at least 4. It can further be seen that if the momentary bit rate at the UNI is less than the negotiated value of NBR, PL is at most 4. The priority level scheme in accordance with this embodiment of the present invention thus permits adjustment of the relative capacity used by a connection in steps of 2. From Equation [1] above, it can be seen that for an NBR of 100 kbit/s, an MBR higher than 566 kbit/sec results in a PL of 7, and an MBR lower than 8.8 kbit/s results in a PL of 0.

The three priority level bits allocated in the cell header are set 78 for each ATM cell transferred from the UNI. The ATM cells are then transmitted 80 to targeted network nodes, identified by node addressing information provided in the cell header.

It is noted that if a user is not satisfied with the quality of service of the connection, the user may elect between at least three alternatives. First, the user may elect to keep the average bit rate

unchanging, but reduce the variation of traffic process. Second, the user may elect to decrease the average bit rate, or to increase the nominal bit rate. Increasing the NBR will, however, generally result in a concomitant
5 increase in cost for a higher speed connection. Finally, the user may change the network operator.

In Fig. 4, there is illustrated in flow diagram form a general methodology by which a network node processes cells containing priority level
10 information received from a UNI in accordance with one embodiment of the present invention. Figure 5 illustrates an embodiment of various components of a network node employed to effectuate the methodology illustrated in Fig. 4. It is assumed that a cell, such
15 as cell_i, has been processed at a UNI and includes priority level information derived in a manner described hereinabove.

Cell_i is transmitted from the UNI to a network node and is received at a filter 88 of the node. A
20 memory manager 89 checks the status 81 of the memory 90 in order to determine the occupancy in the memory 90. The memory manager 89 determines 82 the allowable priority level (PL_a) based on the occupancy state of the memory 90. In general, the memory manager 89
25 establishes a high allowable priority which translates to a low allowable priority "level," for example PL_a = 0 or 2, when the occupancy level of the memory 90 is high (i.e., few available memory locations). When the memory manager 89 determines that the memory 90 has ample
30 capacity for receiving new cells, the memory manager 89 establishes a low allowable priority which translates to a high allowable priority "level," for example PL_a = 6 or 7. As will be appreciated by those skilled in the art, the calculation of PL_a could alternatively be based on
35 unoccupied buffer capacity rather than on buffer

occupancy without departing from the spirit of the invention.

If the priority level of cell_i is greater than the allowable priority level, PL_a, as determined 83 by the memory manager 89, the filter 88 discards 84 cell_i.
5 If, on the other hand, the priority level of cell_i is equal to or less than the allowable priority level PL_a, the filter 88 accepts 85 cell_i. The memory manager 89 coordinates the transfer 86 of cell_i to the memory 90 and
10 updates an index table 91 coupled to the memory manager 89 to include a new index table entry for newly accepted cell_i. In one embodiment, the index table 91 stores the location of the accepted cell_i in the memory 90, and also stores a cell-type indicator which specifies whether
15 cell_i is a real-time cell or a non-real-time cell. As such, the memory 90 may store both real-time and non-real-time cells.

The memory manager 89, in cooperation with the index table 91, manages cell transfer operations from
20 the memory 90 to the output of the memory 90 by giving preference to the real-time cells over the non-real-time cells. By way of example, the memory manager 89, upon determining the presence of both rt-cells and nrt-cells stored in the memory 90, transfers all of the rt-cells
25 to the output of the memory 90 prior to transferring out any of the nrt-cells.

In accordance with another embodiment, as is illustrated in Fig. 6, a memory manager 89 determines the status of a real-time buffer (rt-buffer) 93 and a
30 non-real-time buffer (nrt-buffer) 94. The memory manager 89 determines, in a manner similar to that previously described, the allowable priority level, PL_a, for the filter 88 based on the status of the rt-buffer 93 and the nrt-buffer 94. If the priority level of cell_i
35 is greater than the allowable priority level, PL_a, the

filter 88 discards cell_i. If, on the other hand, the priority level of cell_i is equal to or less than the allowable priority level, PL_a, cell_i is accepted.

In accordance with another embodiment, the network node may apply a buffer filtering scheme which performs the filtering function based on packets of cells, rather than on individual cells. By way of example, the filtering procedure described hereinabove may be applied to the first cell of each packet. If the first cell is discarded by the node, then all of the cells of the packet following the first cell are discarded as well. If, however, the first cell of a packet is accepted, then the priority of all other cells belonging to that packet may be increased, for example by changing the priority level from PL = 5 to PL = 3. A gain of even one priority level, such as from PL = 4 to PL = 3, is believed to be sufficient to ensure that there will only be very few partially transmitted packets.

A cell-type detector 92 receives the accepted cell, cell_i, from the filter 88 and determines whether cell_i is an rt-cell or an nrt-cell. As discussed previously, the header of cell_i includes a header bit, such as the CLP bit, which indicates whether or not cell_i is an rt-cell or an nrt-cell. The cell-type detector 92, upon determining the service class type of the cell_i, transfers the cell_i to either the rt-buffer 93 or the nrt-buffer 94. In a manner similar to that described previously with respect to Figs. 4 and 5, the memory manager 89 coordinates the output of rt-cells and nrt-cells respectively from the rt-buffer 93 and the nrt-buffer 94, giving preference to the rt-cells.

In accordance with another embodiment of the present invention, it may be desirable, for purposes of enhancing network expansion and traffic control, to

request that each user of the network purchase a maximum NBR. The maximum NBR value is intended to remain substantially constant. In addition, it may be desirable to request that each user select an appropriate instantaneous NBR, which should be no greater than the selected maximum NBR. The selection of an appropriate instantaneous NBR generally involves a compromise between price and quality of service. The service quality detected by a user depends largely on three parameters, namely the NBR, the average bit rate, and the amount of traffic variations. Although a user may change any of these parameters, the only information that the network needs to know at the initiation of cell transmission is the NBR and the service class (real-time or non-real-time) of the connection.

In accordance with another embodiment of the present invention, a SIMA service model provides for the accommodation of both NBR and traditional ATM service connections. It is appreciated that traditional ATM services which offer guaranteed connections may be desirable for certain applications. It is anticipated, however, that the quality of service offered by the NBR service of the present invention will meet or exceed a user's expectations for virtually all real-time and non-real-time applications.

A SIMA service which provides for both NBR and traditional ATM services requires that the network operator dedicate a UPC device for each conventional ATM connection, or possibly for each virtual path. All of the cells transmitted using traditional ATM service connections are designated with the highest priority of $PL = 0$ and with a real-time (rt) service class designation. In accordance with this approach, the zero priority level is reserved for those connections that use an ordinary ATM service with guaranteed bandwidth

and quality of service. Accordingly, the priority determination Equation [1] above is modified so as to produce cell priority levels ranging between $PL = 1$ and $PL = 7$, such that $[x]$ represents the integer part of x ,
5 if $1 < x < 7$. It is noted that if the network operator wants to mark excessive cells as $CLP=1$ cells, those cells may be marked with a lower priority, such as $PL = 6$, for example.

A point of possible incompatibility with
10 traditional ATM technology involves the necessity of three bits for each ATM cell for the determination of cell priority, or two bits if the current cell loss priority, CLP, bit in the cell header is used. In addition, one bit is needed to distinguish between real-
15 time and non-real-time connections. The rt/nrt service bit may be, but is not required to be, included in every cell. It may be possible to utilize the current Generic Flow Control (GFC) field which constitutes a total of four bits. In this case, three bits may be allocated
20 for designating cell priority level and one bit may be designated as the rt/nrt service class bit.

Referring now to Fig. 7, one embodiment of a network cell 200 in accordance with the present invention is shown. As previously described, ATM
25 standards define an ATM cell as a fixed-size cell with a length of 53 octets comprised of a 5-octet header and a 48-octet payload. The cell 200 is modeled after the ATM standard cell, and includes a 5-octet header 202 and a 48-octet payload 204. Priority levels are calculated,
30 and a priority level is assigned to each cell prior to being transmitted from a UNI to the network. Designation of the priority level attributed to a particular cell is stored in the priority level field (PL) 206. In the present example, eight priority levels
35 are available, thereby necessitating the allocation of

three bits in the PL field 206.

Fig. 8 is a block diagram illustrating a cell traffic measurement and priority assignment system 250 in accordance with one embodiment of the present invention. The system 250 monitors the cell transmission rate from the UNI on a particular connection, and assigns a priority level designation to each outgoing cell. The priority level assigned to a particular cell depends on the current cell transmission rate as compared to the established NBR. In one embodiment of the invention, eight priority levels exist, and a priority level value indicative of one of the eight priority levels is assigned to each of the cells transferred by a user from the UNI. Assuming eight priority levels, priority level values are assigned in one embodiment such that doubling the cell transfer rate causes the priority level (PL) to be incremented, where a PL of four corresponds to the negotiated nominal bit rate. This is illustrated in Table 1 below:

PL	7	6	5	4	3	2	1	0
Rate	$8 \times \text{NBR}$	$4 \times \text{NBR}$	$2 \times \text{NBR}$	NBR	$\text{NBR}/2$	$\text{NBR}/4$	$\text{NBR}/8$	$\text{NBR}/16$

TABLE 1

25

The cell output source 252 of Fig. 8 is responsible for sending information destined for a particular network node or destination by way of the network 254. As cells are output from the cell output source 252, the virtual cell generator 256 recognizes each of the cells transmitted to the network 254, and generates a stream of corresponding virtual cells at the virtual cell input 258 of the buffer 260. Therefore, the buffer 260 is filled with virtual cells as cells are

30

35

output from the cell output source 252. The current occupancy level of the buffer 252 is indicated by an address which points to the next available buffer location, illustrated as the virtual cell pointer on
5 line 262.

The traffic measurement and priority assignment system 250 is based on the virtual cell buffer 260. The occupancy level of the buffer 260 identifies the priority level which is to be assigned to
10 cells output from the cell output source 252. A high buffer 260 occupancy level generally indicates that the cell output source 252 has been sending cells at a higher rate than the NBR, which results in higher priority level assignments. For example, referring to
15 Table 1, where the traffic measurement and priority assignment system 250 determines that the cell output source 252 is sending cells to the network 254 at a rate twice that of the negotiated NBR (i.e., $2 \times \text{NBR}$), the priority level will be increased to a PL equal to "5".

20 The buffer 260 occupancy level increases by one upon the arrival at the virtual cell input 258 of each cell from the virtual cell generator 256. The buffer 260 is emptied at a rate corresponding to its current occupancy level. In one embodiment of the
25 invention, the cell discharge rate of the buffer 260 is set to approximately the same rate at which the cells arrive at the virtual cell input. In this case, the buffer occupancy level remains at the same level as long as the transmission rate at the cell output source 252
30 does not change. New cells output from the cell output source 252 would be assigned a common priority level based on the unchanging buffer occupancy level.

However, where the cell transfer rate from the cell output source 252 changes such that the buffer 260
35 begins to fill at a different rate, one of a plurality of buffer occupancy thresholds may be crossed. This is detected by the priority threshold compare unit 264,

which determines whether, and which, priority threshold has been crossed according to the buffer occupancy level identified by the cell pointer 262. Where eight priority levels are used, the buffer is divided into
5 eight different occupancy levels, labeled occupancy levels 266, 268, 270, 272, 274, 276, 278 and 280. The occupancy level divisions each correspond to a different priority level. For example, where buffer 260 is filled such that the cell pointer 262 is in the range of
10 occupancy level 268, a priority level value of 1 is presented at the output of the priority threshold compare unit 264, which is entered into the PL field 206 of the cells.

The occupancy levels 266-280 are separated by
15 buffer occupancy thresholds, represented by dashed lines 282, 284, 286, 288, 290, 292, 294 and 296. Each of these predetermined occupancy thresholds are compared to the current occupancy level designated by the cell pointer 262 at the priority threshold compare unit 264
20 to determine to what priority level subsequently transmitted cells should be set. For example, where the buffer occupancy level has increased into occupancy level 274, the priority threshold compare unit 264 recognizes that the threshold represented on line 290
25 has been surpassed, and that a priority value of "4" should be output to the cell output source 252.

The number of virtual cell locations in each of the buffer occupancy levels defines the overall rate of change of adjustment of the priority levels. For
30 example, the larger the range of cells in each occupancy level, the slower the priority level adjustment. The number of buffer locations within the buffer occupancy levels can vary from level to level, as will be described in further detail below.

35 In order to maintain a buffer occupancy level in the buffer 260, the virtual cell output enable circuit 298 inputs the priority value from the priority

threshold compare unit 264, and adjusts the rate at which the buffer 260 discharges the virtual cells at the virtual cell output 300. In one embodiment of the invention, each of the eight priority levels output from the priority threshold compare unit 264 are generally decoded at the cell output enable circuit 298 to provide eight cell transfer rates at the output 300 of buffer 260. For example, the cell output enable circuit 298 can generate one of eight pulse trains, each pulse train having a different frequency which controls the rate at which virtual cells are output from the buffer 260. In another embodiment, the cell pointer value is provided to the cell output enable circuit 298 such that the cell output enable circuit 298 can provide a signal to the output 300 based on the value of the cell pointer value 262.

Referring now to Fig. 9, a flow diagram is shown illustrating one manner of comparing priority threshold values to the current buffer occupancy levels. A first threshold, also the lowest threshold in this example, is compared 300 to the cell pointer value. This may be accomplished by temporarily storing 302 the lowest threshold, such as the threshold represented on line 282, into a register A. The cell pointer is stored 304 into another register B, and registers A and B are compared 306 to determine 308 whether the cell pointer corresponds to an occupancy level greater than the particular threshold value.

If it is determined 308 that the cell pointer value in register B is not greater than the first threshold value in register A, the process continues at block 302. If it is determined 308 that the cell pointer value in register B is greater than the first (e.g., lowest) threshold value in register A, the cell priority is set 310 to the next priority level above the first threshold value. For example, where the first threshold value corresponds to the threshold

corresponding to threshold line 284, the priority level will be set to "1" where the cell pointer is at an occupancy level above line 284. The cell pointer is then compared 312 to the next threshold level in a manner analogous to comparing 300 the cell pointer to the first threshold. If it is determined 314 that the cell pointer value is higher, the cell priority is set 316 to the next priority level above the current threshold. This continues until it is determined 314 that the cell pointer value is not higher than the current threshold, wherein the last cell priority recorded is the priority level value sent from the priority threshold compare unit 264. It should be recognized that the functions illustrated in Fig. 9 are not restricted to a sequence illustrated by the flow diagram, as these functions may in fact be implemented in parallel.

Fig. 10 is a block diagram of a priority threshold compare unit 264 which may perform the comparing function of the priority threshold compare unit 264 in accordance with one embodiment of the present invention. A plurality of comparing units 350, 352, 354, 356, 358, 360, 362 and 364 are used to compare each of the predetermined threshold values represented on lines 282, 284, 286, 288, 290, 292, 294 and 296 to the current cell pointer value 262. The priority threshold compare unit 264 will provide a priority value of "0" on connection 265 via the encoder 368 when compare unit 364 recognizes that the cell pointer is above threshold 282, and all of the other compare units 350-362 indicate that the cell pointer is below their respective thresholds. Both of these conditions are required as indicated by the AND-function 370. The requirement that the other compare units 350-362 indicate a cell pointer value below their respective thresholds is provided through the NOR-function provided via the combination of the OR 372 and INV (inverting)

374 functions.

As will become readily apparent to those skilled in the art, the remaining inputs to the encoder 368 are developed in an analogous manner. Such an arrangement may be implemented using hardware components, or is alternatively conducive to processor-controlled operation.

In one embodiment of the invention, the buffer 260 is implemented in a fashion analogous to that of a counter which can be incremented and decremented, and where the buffer has equally-spaced occupancy levels. The occupancy level of the buffer 260 is examined at each cell arrival at the virtual cell input 258. Therefore, the buffer/counter is not decremented continuously between cell arrivals. Instead, decrementing via the output 300 occurs concurrently with the arrival of a cell at the cell input 258. This implementation simplifies the occupancy level calculations, as calculations need only be performed upon each cell arrival.

Referring now to Fig. 11, a graphical representation of the buffer 260 occupancy level versus the rate at which the buffer 260 is emptied is provided, where the buffer is implemented as a counter. In such a case, the counter values will fall in the range of 0 to $8k$, where k is the number of cells in one buffer occupancy range, such as in occupancy ranges 266-280 shown in Fig. 8. The lower and upper ranges corresponding to priority levels 0 and 7 respectively, illustrated on lines 400 and 402, may have a range wider than one buffer occupancy range k , as shown in Fig. 11. Therefore, where a cell sourcing unit has been providing cells at a rate higher than 8 times the nominal bit rate ($8 \cdot \text{NBR}$) at a priority level of 7, it will take longer to decrease to priority level to 6 than if it had been sending cells at a rate equal to $8 \cdot \text{NBR}$. Correspondingly, if the cell source unit has been

sending cells at a very low rate such that the cells are at a priority level of 0, it will take longer to increase the priority level to a priority level of 1.

In order to simplify the calculations, an approximation of the step function illustrated in Fig. 11 is provided by a continuous line 404. This approximation results in the counter value decrement rate being a continuous function of the buffer occupancy level. This approximation also allows the buffer occupancy level to be represented by a real number rather than an integer value. The approximation is such that it gives the same values as the step function at the points in which the priority value of a cell changes. Therefore, the approximation does not significantly affect the priorities assigned to the cells.

Fig. 12 is a graphical illustration of the buffer occupancy level 420 as a function of time 422 when the buffer 260 is full at time t equals 0, and no cell arrivals occur. This function provides the basis of the approximation of the step function of Fig. 11. In this example the total buffer occupancy level is 80 cells, where k equals 10, and the NBR of the cell sourcing unit is 1,000 cells per second. A high buffer occupancy level indicates that the cell output source 252 has been sending cells at a higher rate than the NBR, which results in a high priority level assignment, which causes the virtual cell output enable circuit 298 to direct a high cell discharge rate from the output 300 of the buffer 260. As the buffer occupancy level decreases over time, the emptying rate at which cells are discharged from the buffer 260 is reduced as a result of a decreasing priority value.

Fig. 13 is a graphical representation of the relationship between the buffer occupancy level, the emptying rate, and various buffer discharge times where the buffer occupancy levels are equally spaced. The

various buffer occupancy level thresholds 282-296 are shown, as are the different occupancy levels 266-280. The corresponding priority level values and buffer occupancy levels are illustrated at rows 440 and 442 respectively. The rate at which the cell output enable circuit 298 enables the output 300 to empty the buffer 260 is illustrated at row 444. The duration of time required to reach an adjacent threshold level from a current threshold level is illustrated at row 446, which illustrates that the time required to reach the next level is an increasing function as the priority level decreases due to the decreasing emptying rate shown on line 444. Lines 448, 450, 452, 454, 456, 458, 460 and 462 represent the times required to reach a particular threshold level starting from the time corresponding to a full buffer 260. For example, line 450 represents the time that it would take for a full buffer 260 to decrease to an occupancy level corresponding to the predetermined threshold value of line 294 where no cell arrivals occur. Because the cell discharge rate decreases by a factor of 2 as the priority level decreases, the time required to reach the next threshold level doubles. Therefore, the time on line 450 is derived by adding the time corresponding to the occupancy level 278 to the time corresponding to occupancy level 280 from line 448. The times required to reach particular levels from a full buffer was depicted graphically in Fig. 12.

By understanding the nature in which the buffer is emptied, the number of cells that have been output from the buffer 260 since a previous cell arrival can be calculated. In one embodiment of the invention, this is accomplished by calculating the time (t) taken to empty a full buffer 260 to the current occupancy level (OL), adding the time passage from the arrival of the previous cell to the arrival time of the currently-received cell (dt), and evaluating the buffer occupancy

level at the current time. The current cell arrival is then added to the calculated buffer occupancy level to determine the total occupancy level at the time of the arrival of the current cell.

5 Fig. 14 graphically illustrates the calculation of the number of cells discharged from the buffer 260 between any two cell arrivals. The time (t) corresponding to the time required to empty a full
10 buffer to the existing occupancy level (i.e., the occupancy level after arrival of the previous cell) is calculated as shown in Equation 2 below:

$$t = \frac{k}{8 \cdot NBR} \cdot [2^{8-OL/k} - 1] \quad [2]$$

15

The occupancy level (OL) at the cell arrival is determined at time (t+dt) using Equation 3 below:

$$20 \quad OL(t+dt) = \left[8 - \log_2 \left(\frac{NBR}{k} \right) \cdot 8 \cdot (t+dt) + 1 \right] \cdot k + 1 \quad [3]$$

Equation 3 above determines the occupancy level at the current cell arrival time, taking into account the additional one cell which arrived as well.
25 The occupancy level is converted to a corresponding priority level in the manner described in Equation 4 below:

$$PL = \min \{ 7, \max [0, \text{truncate}(OL/k)] \} \quad [4]$$

30

Referring again to Fig. 14, the time (t) represented on line 500 corresponds to the required to empty a full buffer to the current occupancy level prior to receiving the new cell, which can be determined using

Equation 2 above. The time (dt) represented on line 502 corresponds to the time from the previous cell arrival until the current cell arrival. The total time (t+dt) illustrated on line 504 therefore corresponds to the
5 time elapsed after the new cell arrival. The occupancy level (OL) can then be calculated using Equation 3 above, and the priority level (PL) is determined using the occupancy level calculated in Equation 4.

Fig. 15 is a flow diagram illustrating one
10 embodiment of the invention for the equally spaced buffer occupancy level implementation described in connection with Figs. 11-14. When a new cell arrives 550, the time of arrival is set 552 to equal to the current time. Because the time of the previously
15 received cell arrival is known, the time period between the arrival of the previous cell and the current cell can be calculated 554 as the time of arrival minus the time of the previous cell arrival. This new cell arrival time is then set 556 to the time of previous
20 arrival for use when the next cell arrives. The elapsed time is calculated 558 by adding the calculated time between cells 554 and the time corresponding to the current occupancy level of the buffer 260, which is calculated using Equation 2. The occupancy level can
25 then be calculated 560 Equation 3. The total occupancy level is calculated 562 by adding the single cell arrival to the occupancy level previously calculated 560. The occupancy level is then set 564 to the maximum of the lower margin, and a value corresponding to the
30 minimum of either the occupancy level or the upper margin plus 8k. The priority level is calculated 566 using the calculated occupancy level 564, as was shown in Equation 4. The calculated priority value is assigned 568 to cells being output from the cell output
35 source 252.

In another embodiment of the invention, the buffer 260 is again implemented analogously to that of a

counter, but where the buffer has logarithmically-spaced buffer occupancy levels. With equally-spaced buffer occupancy levels, the time required to decrease one priority level depends on the current priority level.

5 For example, the higher the priority value, the faster the priority value is decremented, as was illustrated in Fig. 12. This condition, however, can be modified by adjusting the buffer occupancy levels that establish the priority values. In one embodiment of the invention,
10 the occupancy levels are set such that the time taken to decrease one priority level is the same, regardless of the current occupancy level. This is illustrated in Fig. 16 below, where each occupancy range includes twice as many cells as the previous occupancy range.

15 Fig. 16 is a graphical representation of the relationship between the buffer occupancy level, the emptying rate, and various buffer discharge times where the buffer occupancy levels are logarithmically spaced. The various buffer occupancy level thresholds 282-296
20 are shown, as are the different occupancy levels 266-280. The corresponding priority level values, buffer occupancy ranges, and buffer occupancy levels are illustrated at rows 440, 600 and 442 respectively. As seen from Fig. 16, the buffer occupancy level on line
25 442 is doubled for each increasing occupancy range 266-280. The buffer occupancy ranges shown on line 600 illustrate that the occupancy range doubles at the priority level increases. The rate at which the cell output enable circuit 298 enables the output 300 to
30 empty the buffer 260 is illustrated at row 444. The time required to reach an adjacent threshold level from a current threshold level is illustrated at row 446, which illustrates that the time required to reach the next level is not an increasing function as the priority
35 level decreases as was illustrated in Fig. 13, but rather is an equal time. Lines 602, 604, 606, 608, 610, 612, 614 and 616 represent the times required to reach a

particular threshold level starting from the time corresponding to a full buffer 260. As can be seen, the time to reach these various levels are determined simply by taking the sum of each of the equal times t_0 shown on line 446 until reaching the appropriate level. For example, line 608 represents the time that it would take for a full buffer 260 to decrease to an occupancy level corresponding to the predetermined threshold value of line 290 where no cell arrivals occur. The time on line 608 is derived by adding the time corresponding to the occupancy levels 280, 278 and 276 to the time corresponding to occupancy level 274, all of which are an equal time of t_0 .

Fig. 17 is a graphical illustration of the buffer occupancy level 650 as a function of time 652 when the buffer 260 is full at time t equals 0, and no cell arrivals occur where the buffer occupancy levels are logarithmically spaced. The priority transition points 654, 656, 658, 660, 662, 664, 666, 668 and 670 designate the decreasing occupancy level as a function of time. The time between each of the priority transition points 654-670 is equal, while the number of cells discharged during each increasing time interval is reduced by a factor of two.

The time (t) corresponding to the current occupancy level for the buffer discharge times where the buffer occupancy levels are logarithmically spaced is calculated as shown in Equation 5 below:

$$t = \frac{16k}{NBR} [8 - \log_2(OL/k + 1)] \quad [5]$$

The occupancy level (OL) at the time of cell arrival is determined using Equation 6 below, where time dt is the elapsed time from the arrival of the previous cell:

$$OL(t+dt) = (2^{8-t/t_0} - 1) \cdot k + 1 \quad [6]$$

$$\text{where } t_0 = 16 \left(\frac{k}{NBR} \right).$$

The priority value assigned to the cell is calculated
5 using Equation 7 below:

$$PL = \log_2 \left(\frac{OL}{k} + 1 \right) \quad [7]$$

Fig. 18 illustrates in flow diagram form a
10 general methodology for which the output rate of the
virtual cells from the buffer 260 is modified. The
priority level from the priority threshold comparator
264 is read 700, and the corresponding value is decoded
702 into n levels, where n depends on the number of
15 priority levels implemented. It should be recognized
that the cell pointer value can alternatively be
directly read, rather than the priority value from the
priority threshold comparator 264, where the cell
pointer value would be directly used to allow the cell
20 output enable circuit 298 to control the output 300.

As previously indicated, one embodiment of the
invention includes eight priority levels which can be
encoded into three bits, resulting in a three bit
priority level and 8 decoded levels. One of n output
25 enable signals is generated 704, which in turn controls
the rate at which virtual cells are discharged 706 from
the buffer 260. When a priority level change is
detected 708, the output enable signal is reconfigured
according to the new priority level. Otherwise, until
30 the priority level changes 708, the virtual cells
continue to be output 706 at the same rate.

In another embodiment of the invention, the
cell pointer value is read by the cell output enable
circuit 298, and one of the n output enable signals is

generated therefrom by analyzing the cell pointer value rather than receiving a priority value from the priority threshold compare unit 264. In such a case, the virtual cell output enable circuit 298 must determine which
5 buffer occupancy level the cell pointer is associated with. This can be accomplished using hardware components in a manner analogous to the previous description of the priority threshold compare unit 264, or can alternatively be accomplished using known
10 software techniques.

Referring now to Fig. 19, one example of a cell output enable circuit 750 is illustrated. An oscillating input clock signal shown on line 752 is input into a frequency division unit 754. The frequency
15 division unit 754 provides a plurality of output signals of various multiples of the clock frequency of the clock signal on line 752. In this example, an 8:1 frequency division occurs resulting in eight different output enable signals on lines 756, 758, 760, 762, 764, 766,
20 768 and 770. The generated signal waveforms are fed into a multiplexing unit 772. The multiplexing unit 772 is controlled by control lines 774, 776 and 778, which, in one embodiment of the invention, comprises the priority value on connection 265 shown in Fig. 8.
25 Depending on the state of the control lines, the multiplexing unit 772 selects one of its inputs on line 756-770 to be output to the output circuit 300 on line 780. The output circuit 300 allows virtual cells from the buffer 260 to be discharged at each occurrence of a
30 pulse or transition level of the signal on line 780.

Other embodiments of cell output enable circuits 298 and 750 can be implemented without departing from the scope and spirit of the invention. For example, the signals on control lines 774, 776 and
35 778 can be input into a 3:8 decoder, where each of the eight decoded signals feeds an individual tri-state driver whose outputs are commonly coupled as the output

enable signal. Each of the tri-state drivers would have a different input clock signal which would be selected when its respective enable input was activated by a selected signal on the 3:8 decoder. In yet another
5 embodiment, the cell output enable circuit 298 can be a computing device, or a portion thereof, operating under the control of software that regulates the rate at which virtual cells are output from the buffer 260.

For purposes of illustration, and not of
10 limitation, examples are provided below to illustrate the relationship between the quality of service of an NBR or SIMA connection and throughput at different priority levels. The following examples demonstrate, in
15 accordance with one embodiment of the present invention, the relative differences in QoS with respect to adjacent priorities, such as the QoS associated with $PL = 4$ in comparison to that of $PL = 3$. It is noted that a higher cost or fee is typically assessed to users who request a
20 higher priority for cells transmitted from the user's user/network interface. By way of example, the user fee may be doubled if the user wants to obtain one degree of higher priority for every cell without changing the
25 actual bit rate. The resulting QoS of the connection, therefore, should be improved such that at least some users are willing to pay the additional charge.

EXAMPLE #1

In accordance with this illustrative example, the following assumptions and consideration are given.
30 It is assumed that there are many identical traffic sources which generate traffic independent of the current or previous load conditions in the network. The following traffic parameters are assumed: the link capacity is $C = 1$, which is useful in the examples as a
35 means of normalization; the peak bit rate $MBR_{MAX} = 0.1$,

which represents 10 percent of the link capacity, C ; the ON probability at the burst (or packet) scale = 0.2; and the average burst duration = 1,000 time slots (i.e., the average packet size = 100 cells). In addition, it is
5 assumed that there is an upper ON/OFF layer, and that both the average ON-period and OFF-period of this layer are 100,000 time slots. The real time buffer 93 contains 200 cell locations and the non-real-time buffer 94 contains 5,000 cell locations. It is noted that the
10 upper ON/OFF layer attempts to model the traffic process of connections, in which the determination of the number of connections is understood in the art to constitute a stochastic, random process. For example, if it is assumed that the total number of customers is
15 represented by the variable x , then the average number of connections is $x/2$. More particularly, the number of connections is understood to be binomially distributed. As such, 100,000 time slots represent the average holding time of a connection, and, also, the average
20 idle period realizable by the user. As a consequence, a user is transmitting cells only if a connection is active both at the connection layer and the packet layer. A time scale parameter, α , can be obtained for the real-time and non-real-time connections:

25

$$\alpha_{rt} = 0.025$$

$$\alpha_{nrt} = 0.001$$

In this example, eight different connection
30 types are assumed: four connections are real-time connections and four are non-real-time connections. Also, four different NBR values, which have been normalized with respect the link capacity of $C = 1$, are assumed as: 0.2, 0.1, 0.05 and 0.025. The priorities

3, 4, 5 and 6, respectively. It should be noted, however, that not all cells are assigned these exact priorities, and that especially with non-real-time connections, many cells obtain better priority values because of the affects of the averaging measuring principle. The distribution of cells having different priority levels, represented as percentages, is presented below in Table 2:

10

PRIORITY LEVEL	REAL (SIMULATED) PERCENTAGE OF OFFERED CELLS	PERCENTAGE BASED ON PEAK RATES
1	6.1	0
2	7.9	0
3	24.3	25
4	23.5	25
5	21.5	25
6	16.8	25

TABLE 2

In Fig. 20, there is shown a graph illustrating the relationship of average cell loss ratio, P_{loss} , as a function of priority level for four specific load levels, r . In particular, line-800 represents an overall average load level of 0.72 for 9 connections of each connection type (i.e., real-time and non-real-time connection types). Line-802 depicts an average load level of 0.80 for 10 connections of each connection type. Further, line-804 represents an average load level of 0.88 for 11 connections of each connection type, and line-806 represents an average load level of 0.96 for 12 connections of each connection type. It is noted that, in the case of line-802 indicating a load level of 0.80, the cell loss ratios, P_{loss} , for real-time and non-real-time cells are indicated by dotted and broken lines, respectively.

Given, for example, a traffic scenario where the operator wants to offer a cell loss ratio of 10^{-6} to cells with priority 4, the total load can be approximately 0.75. It can be assumed that this average
5 cell loss ratio is sufficient for most video applications. Give the same traffic load conditions, priority level 5, which corresponds to $P_{\text{loss}} \approx 10^{-4}$, can meet the requirements of many voice applications, while priority 6, which corresponds to $P_{\text{loss}} \approx 3 \cdot 10^{-3}$, is
10 suitable for a TCP/IP type of file transfer, provided that there is an adequate packet discarding scheme in place.

It should be emphasized, however, that the difference in cell loss ratio between adjacent
15 priorities depends strongly on the offered traffic process and, in particular, the inherent control loops of the NBR or SIMA service. When the user perceives an unsatisfactory QoS, for example, the user can, and should, change either the actual bit rate or the nominal
20 bit rate of the connection. In either case, the priority distribution changes as well. Nevertheless, if this phenomenon is temporarily ignored, the basic behavior of priority distribution may be further appreciated by making the following simplifying
25 assumption: If it is assumed that all traffic variations are slow as compared to the measuring period and buffer size, then a well-known, conventional ATM approach to approximating cell loss ratio may be used, with the additional requirement that the eight NBR
30 priority levels are taken into account.

If the loss ratio of cells with priority k is denoted by $P_{\text{loss},k}$, and the average loss ratio of cells with a priority of 0 to k is denoted by $P_{\text{loss},k}^*$, then the following equation, which ignores buffering effect,
35 provides that:

$$P_{\text{loss},k}^* = \frac{\sum_{j: \lambda_j > c} \Pr\{\lambda_k^* = \lambda_j\} (\lambda_j - c)}{\rho_k^* c}$$

$$P_{\text{loss},0} = P_{\text{loss},0}^*$$

5

$$P_{\text{loss},0} = \frac{\rho_k^* P_{\text{loss},k} - \rho_{k-1}^* P_{\text{loss},k-1}^*}{\rho_k^* - \rho_{k-1}^*} \quad \text{for } k = 1 \dots 7$$

[8]

where, λ_k^* represents the momentary bit rate level of all
 10 cells with a priority of 0 to k, ρ_k^* represents the
 average offered load produced by these cells, and c
 represents the link capacity. The probability $\Pr\{\lambda_k^* = \lambda_j\}$
 can be calculated in a straightforward manner using
 known convolution techniques.

15

EXAMPLE #2

For purposes of further illustration, a second
 example is provided which assumes the same sources
 described in Example #1, except for the long ON and OFF
 20 periods. Because of the long periods reflected in
 Example #1, the peak rate always determines the cell
 priority. As the buffers are typically not capable of
 filtering any traffic variations, the allowed load in
 Example #2 is much lower than that in the original case
 25 of Example #1.

In Fig. 21, there is illustrated in graphical
 form a relationship between cell loss ratio as a
 function of priority level for different load levels, r.
 It is assumed in Fig. 21 that the peak cell rate of each
 30 connections depicted by solid lines 820, 822, 824 is

0.1, the peak cell rate of each connection depicted by the broken line-826 is 0.2, and the peak cell rate of each connection depicted by the dotted line-828 is 0.05.

Figure 21 shows the cell loss probabilities obtained by application of Equation [8] for different priorities, depicted by the three solid lines, line-820, line-822, and line-824. In addition, two slightly different traffic cases are represented by the dotted line-828 and the broken line-826. The effects of changing the traffic variations are reflected in the graph provided in Fig. 21. The actual change in traffic variations is a direct consequence of doubling or halving bit rates and NBR values.

In a network that embraces the NBR/SIMA service concept, an increase of traffic variations has two main effects if the operator keeps the QoS of priority level 4 unchanged. First, the allowed load level is decreased in the same way as in conventional ATM, and second, the difference in cell loss ratio between adjacent priority level decreases. For purposes of providing a rough estimate of QoS based on Figs. 20 and 21, it may be assumed that if priority level 4 offers a cell loss probability of 10^{-6} , then the cell loss probability will be approximately 10^{-4} to 10^{-3} with priority level 5 depending on the overall traffic variations. The cell loss ratio with priority level 3 may be assumed to be less than 10^{-9} , unless the traffic variations are very pronounced.

Although the above examples provide illustrations of relationships between QoS and priority levels, it may be unfruitful to attempt to exactly determine the allowed load or the cell loss difference between adjacent priority levels until user reactions to different QoS and usage charges are evaluated. In an NBR/SIMA service environment, a schedule of charges

based on different QoS levels may be determined, in a certain sense, automatically. For example, if the difference in cell loss ratio between priority levels 4 and 5 is very small, it can be assumed that some of the connections will tend to move from priority level 4 to level 5 because of a lower assessed charge. This change indicates, apparently, that the cell loss ratio of priority level 4 decreases and the cell loss ratio of priority level 5 increases. It can be reasonably assumed that this type of movement continues until the QoS difference corresponds to the average user's expectation of a reasonable charging structure.

Similar concerns are raised with regard to the differences in charging which occur automatically during busy hours in contrast to idle hours. For example, it would appear reasonable to charge higher prices during periods of high load for a certain QoS, and charge lower prices during low load periods. However, it is believed desirable to avoid instigating a charging policy that assesses different fees for a certain NBR during busy and idle periods, which also avoids increasing the complexity of the charging scheme. The naturally occurring "supply and demand" effect may tend to automatically even out the load between busy and idle hours. It is anticipated that, if a user is displeased with the observable difference in QoS during busy and idle periods, the user may be motivated to pay a different rate during such periods.

It will, of course, be understood that various modifications and additions can be made to the preferred embodiments discussed hereinabove without departing from the scope of the present invention. Accordingly, the scope of the present invention should not be limited by the particular embodiments discussed above, but should be defined only by the claims set forth below and

equivalents thereof.

CLAIMS

What is claimed is:

5

1. A method of adjusting a transmission rate of information cells transmitted over a network connection, comprising:

transmitting a measurement cell to a
10 measurement buffer substantially contemporaneously with respect to each transmission of an information cell to the network connection;

regulating an occupancy level of the measurement buffer;

15 generating a cell priority value in response to the occupancy level of the measurement buffer, the cell priority value indicating a probability of information cell acceptance at a node of the network connection; and

20 assigning the cell priority value to information cells being transmitted to the network connection.

2. The method of claim 1, wherein generating the
25 cell priority value further comprises determining the occupancy level of the measurement buffer.

3. The method of claim 1, wherein regulating the occupancy level of the measurement buffer further
30 comprises increasing a discharge rate of the measurement buffer in response to an increase in the occupancy level of the measurement buffer and decreasing the discharge rate of the measurement buffer in response to a decrease in the occupancy level of the measurement buffer.

35

4. The method of claim 1, wherein generating the cell priority value further comprises decreasing the cell

priority value in response to an increase in the occupancy level of the measurement buffer and increasing the cell priority value in response to a decrease in the occupancy level of the measurement buffer.

5

5. The method of claim 1, wherein the measurement buffer is arranged to include a plurality of buffer occupancy ranges, and generating the cell priority value further comprises generating the cell priority value that varies relative to the occupancy level falling within a particular one of the plurality of buffer occupancy ranges.

6. The method of claim 5, wherein regulating the occupancy level of the measurement buffer further comprises increasing a discharge rate of the measurement buffer in response to the occupancy level increasing from a current buffer occupancy range to a higher buffer occupancy range, and decreasing the discharge rate of the measurement buffer in response to the occupancy level decreasing from the current buffer occupancy range to a lower buffer occupancy range.

7. The method of claim 5, wherein generating the cell priority value further comprises decreasing the cell priority value in response to the occupancy level increasing from a current buffer occupancy range to a higher buffer occupancy range, and increasing the cell priority value in response to the occupancy level decreasing from the current buffer occupancy range to a lower buffer occupancy range.

8. The method of claim 1, wherein the network connection comprises a nominal bit rate connection.

35

9. A method of claim 1, and further comprising:

configuring the measurement buffer to include a plurality of buffer occupancy ranges;

discharging measurement cells from the measurement buffer at a discharge rate that varies
5 relative to a current occupancy level of the measurement buffer falling within a particular range of the plurality of buffer occupancy ranges; and

generating a cell priority value having a magnitude that varies relative to the current occupancy
10 level of the measurement buffer falling within the particular buffer occupancy range.

10. The method of claim 9, wherein the magnitude of the cell priority value decreases in response to a
15 decrease in the discharge rate of measurement cells from the measurement buffer, and the magnitude of the cell priority value increases in response to an increase in the discharge rate of measurement cells from the measurement buffer.

20

11. The method of claim 9, wherein generating the cell priority value further comprises decreasing the magnitude of the cell priority value in response to the occupancy level of the measurement buffer increasing
25 from the particular buffer occupancy range to a higher buffer occupancy range, and increasing the magnitude of the cell priority value in response to the occupancy level of the measurement buffer decreasing from the particular buffer occupancy range to a lower buffer
30 occupancy range.

12. The method of claim 9, wherein the time to discharge the measurement cells from the measurement buffer remains substantially constant in a change from
35 one cell priority value to the next.

13. The method of claim 9, wherein each of the buffer occupancy ranges includes substantially the same number of cell storage locations, and a duration of time required to fill each of the buffer occupancy ranges
5 increases exponentially as the current buffer occupancy level increases from the particular buffer occupancy range to a higher buffer occupancy range.

14. The method of claim 9, wherein each of the
10 buffer occupancy ranges includes an exponentially larger number of cell storage locations than a preceding buffer occupancy range, and a duration of time required to fill each of the buffer occupancy ranges is substantially the same as the time required to fill other ones of the
15 buffer occupancy ranges.

15. A system for adjusting a transmission rate of information cells transmitted over a network connection, comprising:
20 a source unit that transmits the information cells to the network connection;
a measurement buffer that receives measurement cells from the source unit, the source unit transmitting a measurement cell to the measurement buffer
25 substantially contemporaneously with each information cell transmission;
a detector that detects an occupancy level of the measurement buffer, the detector producing an occupancy level signal indicative of a current occupancy
30 level of the measurement buffer;
an output enable circuit, coupled to the detector, that controls a discharge rate of measurement cells from the measurement buffer in response to the occupancy level signal; and
35 an encoder that encodes a priority value in each of the information cells, the priority value having

a magnitude that varies in response to the current occupancy level of the measurement buffer.

16. The system of claim 15, wherein the detector
5 produces the occupancy level signal that varies relative to the current occupancy level of the measurement buffer.

17. The system of claim 15, wherein the
10 measurement buffer is arranged to include a plurality of buffer occupancy ranges demarcated by a corresponding plurality of threshold values, and the detector produces the occupancy level signal that varies relative to detection of a particular one of the plurality of
15 threshold values.

18. The system of claim 15, wherein the output enable circuit increases the discharge rate of measurement cells from the measurement buffer in
20 response to an increase in the occupancy level of the measurement buffer, and decreases the discharge rate of measurement cells from the measurement buffer in response to a decrease in the occupancy level of the measurement buffer.

25 19. The system of claim 15, wherein the encoder encodes a priority value of a decreasing magnitude in response to an increase in the occupancy level of the measurement buffer, and encodes a priority value of an
30 increasing magnitude in response to a decrease in the occupancy level of the measurement buffer.

20. A method for controlling a cell priority for cells transmitted from a network source unit to one or
35 more network destination units across a network connection, comprising:

registering each of the cells transmitted from ..
the network source unit in a measurement buffer;

discharging the registered cells in the
measurement buffer at a discharge rate corresponding to
5 a current buffer occupancy level of the measurement
buffer;

generating a cell priority value correlating
to one of a plurality of predetermined buffer occupancy
ranges which coincides with the current buffer occupancy
10 level; and

loading the cell priority value into the cells
being transmitted from the network source unit.

21. The method of Claim 20, further comprising
15 establishing a nominal bit rate for the cells
transmitted from the network source unit, and for
presetting the cell priority value to a midpoint of a
continuum of the priority values.

20 22. The method of Claim 20, wherein registering
each of the cells comprises providing a virtual cell in
the measurement buffer for each of the cells transmitted
from the network source unit.

25 23. The method of Claim 22, wherein providing a
virtual cell comprises generating a duplicate cell for
each of the cells transmitted from the network source
unit and storing the duplicate cell in the measurement
buffer.

30 24. The method of Claim 22, wherein providing a
virtual cell comprises incrementing a buffer address
pointer by one cell amount for each of the cells
transmitted from the network source unit.

35

25. The method of Claim 24, wherein discharging the registered cells comprises decrementing the buffer address pointer according to the discharge rate.

- 5 26. The method of Claim 20, wherein registering each of the cells comprises buffering the cells transmitted from the network source unit in the measurement buffer.
- 10 27. The method of Claim 20, wherein generating the cell priority value comprises identifying the buffer occupancy range that coincides with the current buffer occupancy level by comparing the current buffer occupancy level to a plurality of threshold values
- 15 separating the plurality of predetermined buffer occupancy ranges.

28. The method of Claim 20, wherein generating a cell priority value comprises encoding the cell priority value into n bits, wherein a number of the predetermined buffer occupancy ranges is less than or equal to 2^n .
- 20

29. The method of Claim 20, wherein discharging the registered cells comprises regulating the discharge rate by providing one of a plurality of output enable signals to a measurement buffer output, wherein a different one of the plurality of output enable signals is provided for each of the predetermined buffer occupancy ranges.
- 25

- 30 30. The method of Claim 20, wherein discharging the registered cells comprises regulating the discharge rate by providing an output enable signal proportional to the current buffer occupancy level to a measurement
- 35 buffer output.

31. The method of Claim 20, wherein discharging comprises increasing the discharge rate as the current occupancy level increases from a current buffer occupancy range to a higher one of the buffer occupancy ranges, and decreasing the discharge rate as the current occupancy level decreases from the current buffer occupancy range to a lower one of the buffer occupancy ranges.

10 32. The method of Claim 20, wherein the each of the predetermined buffer occupancy ranges includes substantially the same number of cell locations, such that the time required to fill each of the buffer occupancy ranges increases exponentially as the current
15 buffer occupancy range, and the cell priority value, increase.

33. The method of Claim 20, wherein each of the predetermined buffer occupancy ranges includes an
20 exponentially larger number of cell locations than a preceding buffer occupancy range, such that the time required to fill each of the buffer occupancy ranges is substantially the same as the time required to fill other ones of the buffer occupancy ranges.

25

34. A system for adjusting a transmission rate of information cells transmitted over a network connection, comprising:

30 a source unit that transmits the information cells to the network connection;

a cell transmission measurement device coupled to the source unit to register an occurrence of each information cell transmission to provide a quantitative cell transmission level;

35 a detector that detects the cell transmission level of the cell transmission measurement device, the detector producing a cell level signal indicative of a

current quantitative cell transmission level of the cell transmission measurement device;

an output enable circuit, coupled to the detector to control a decrementing rate of the quantitative cell transmission level from the cell transmission measurement device in response to the cell level signal; and

an encoder that encodes a priority value in each of the information cells, the priority value having a magnitude that varies in response to the current quantitative cell transmission level of the cell transmission measurement device.

35. The system of Claim 34, wherein the cell transmission measurement device is a buffer capable of storing a virtual cell upon the occurrence of each information cell transmission, and wherein an occupancy level of the buffer corresponds to the quantitative cell transmission level.

36. The system of Claim 34, wherein the cell transmission measurement device and the output enable circuit collectively comprise a counter, wherein the quantitative cell transmission level corresponds to a counter value of the counter, and wherein the counter value is incremented upon the occurrence of each information cell transmission, and is decremented at a rate dictated by the cell level signal.

37. The system of Claim 34, wherein the decrement rate of the counter value increases as the counter value increases.

38. The system of Claim 34, wherein the encoder encodes a priority value in each of the information cells, wherein the priority value has a magnitude

corresponding one of a predetermined plurality of ranges..
of the counter values.

5

10

1/18

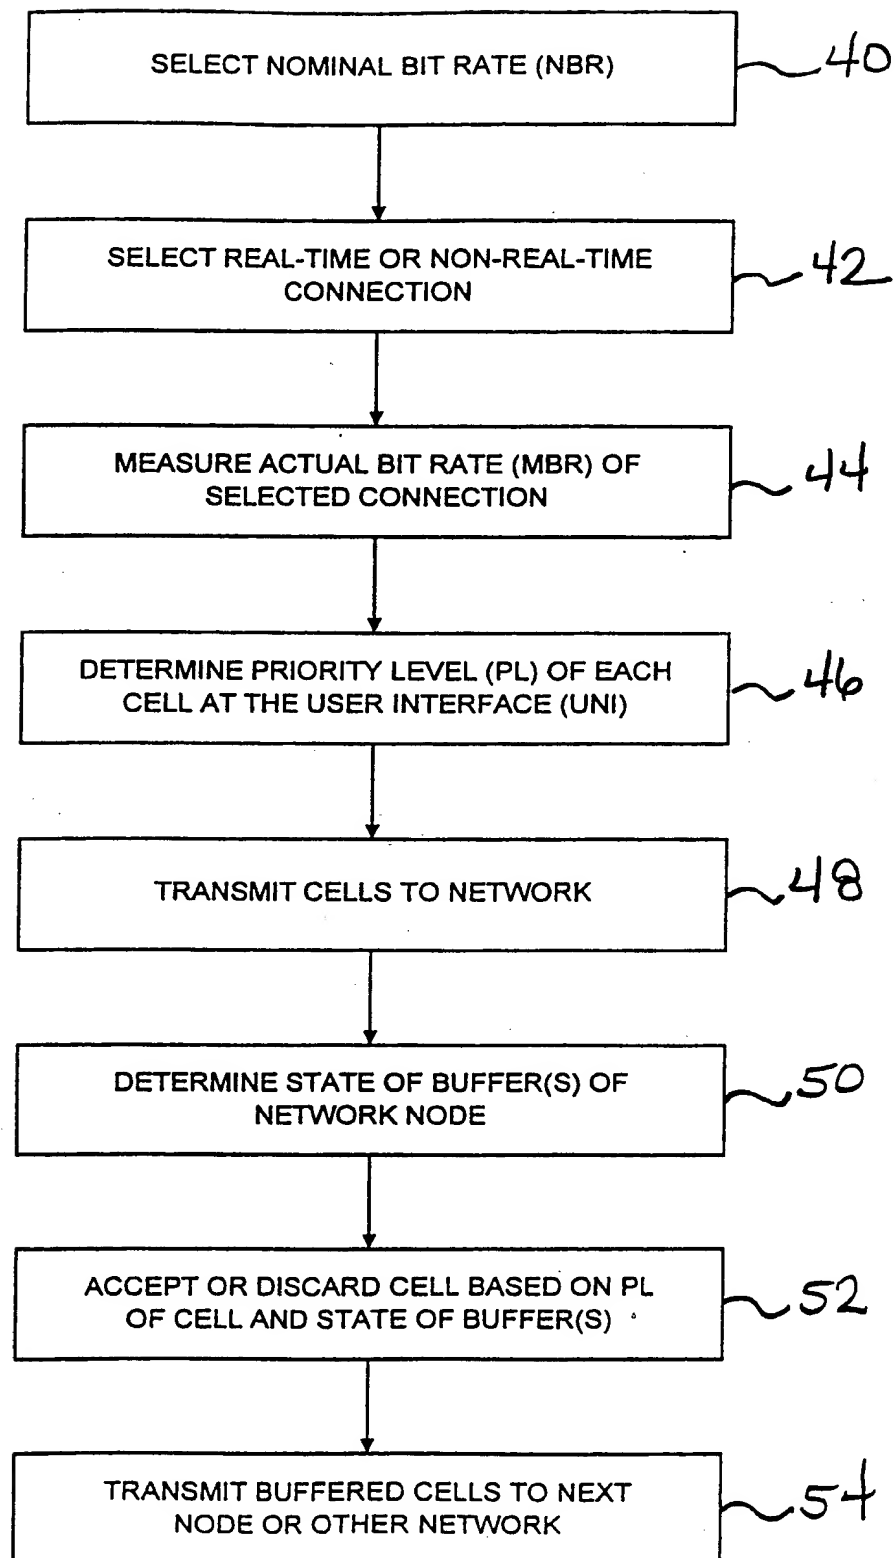


FIG. 1

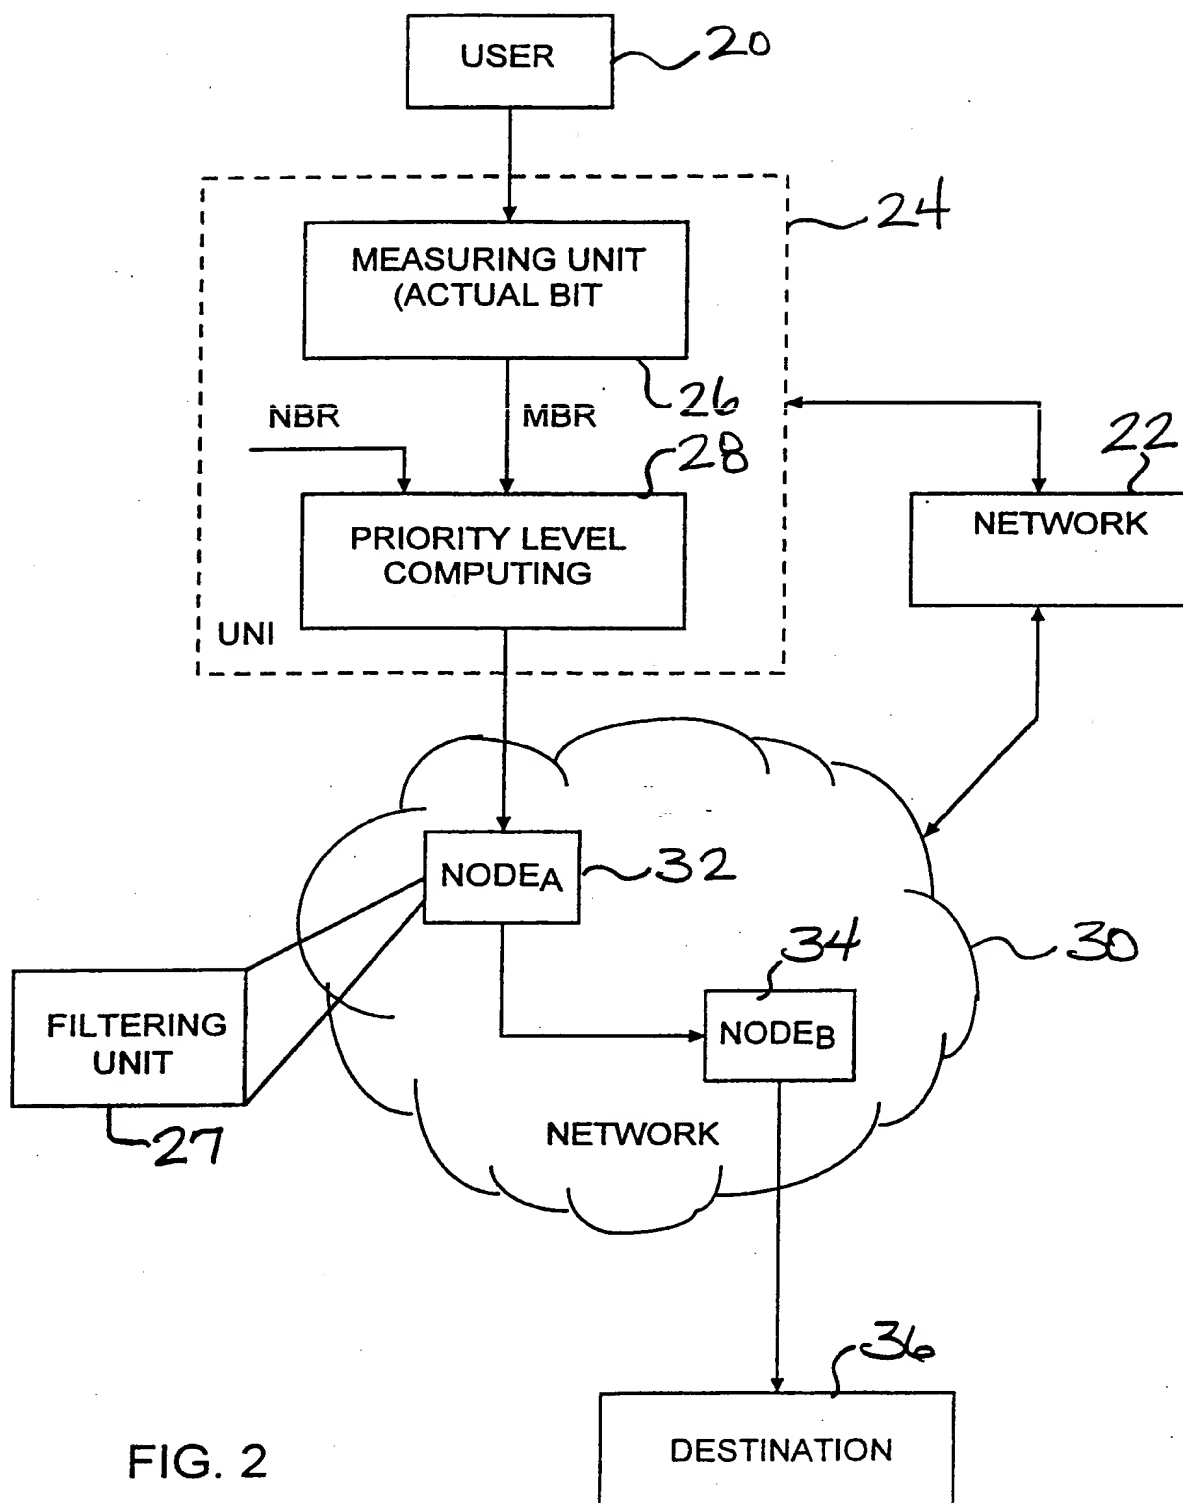


FIG. 2

3/18

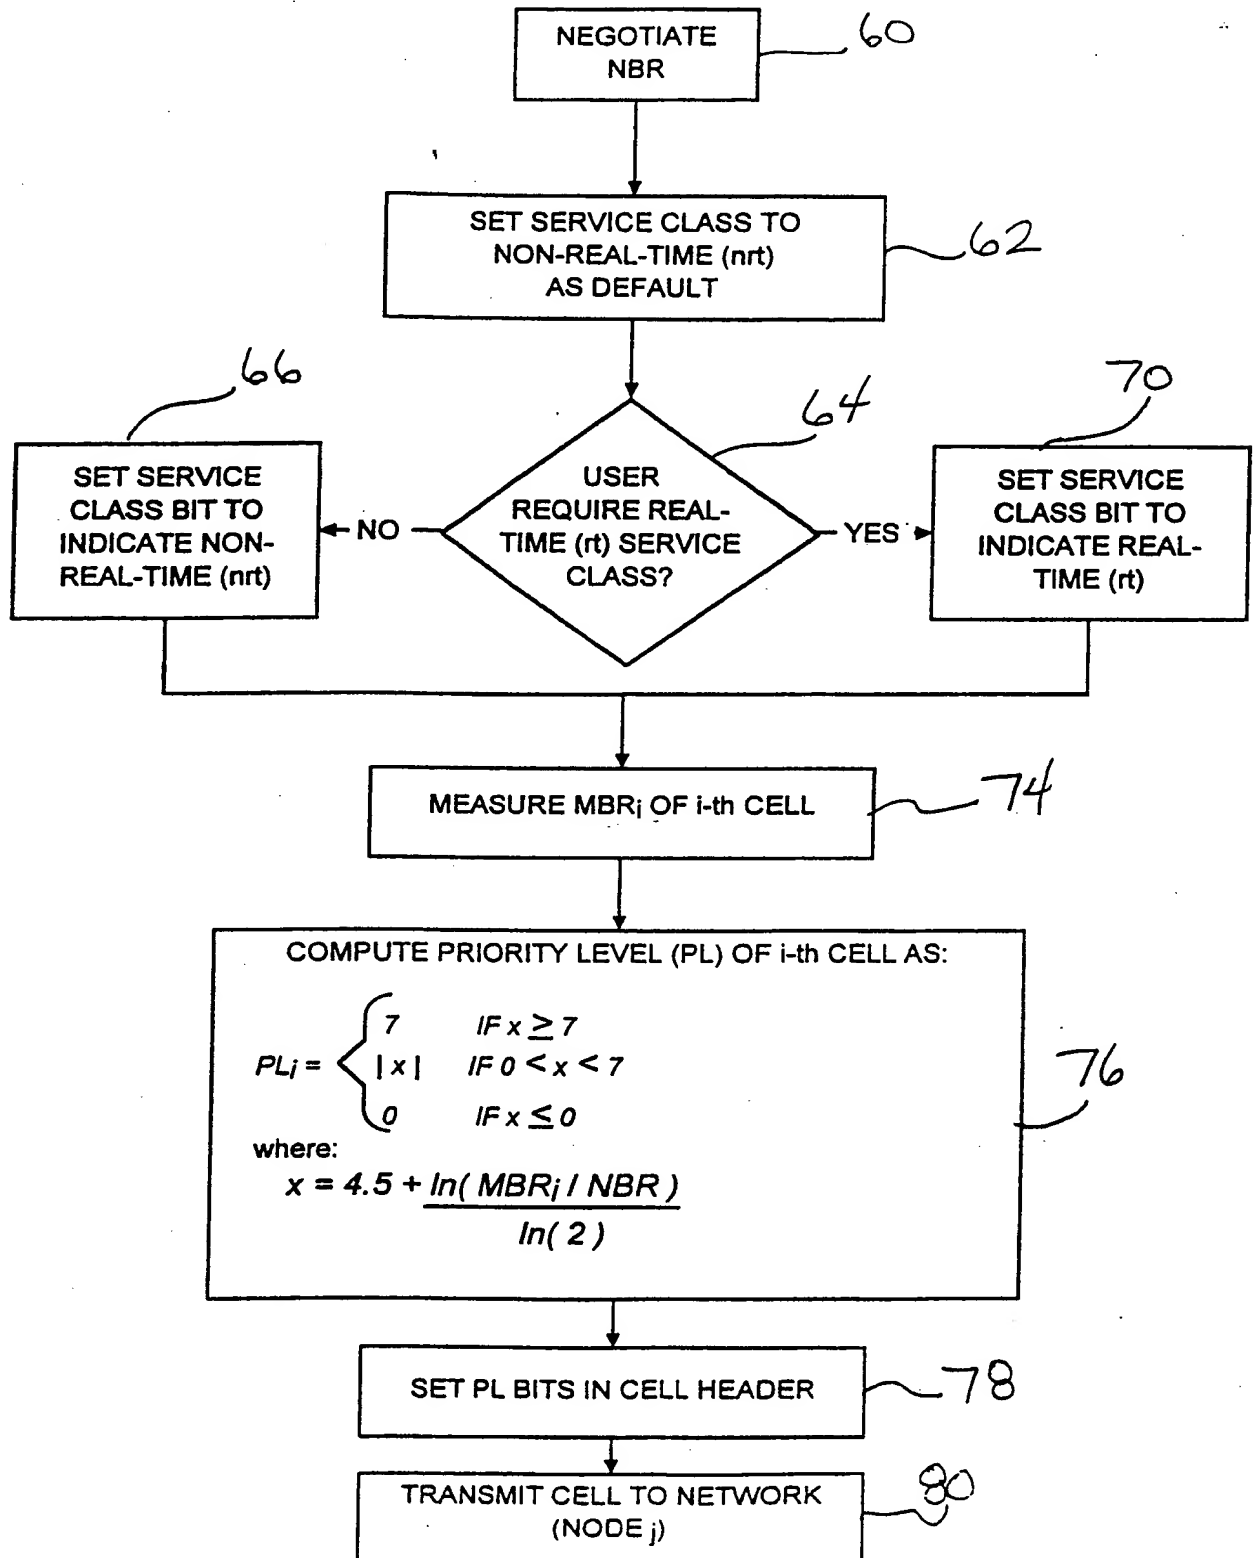
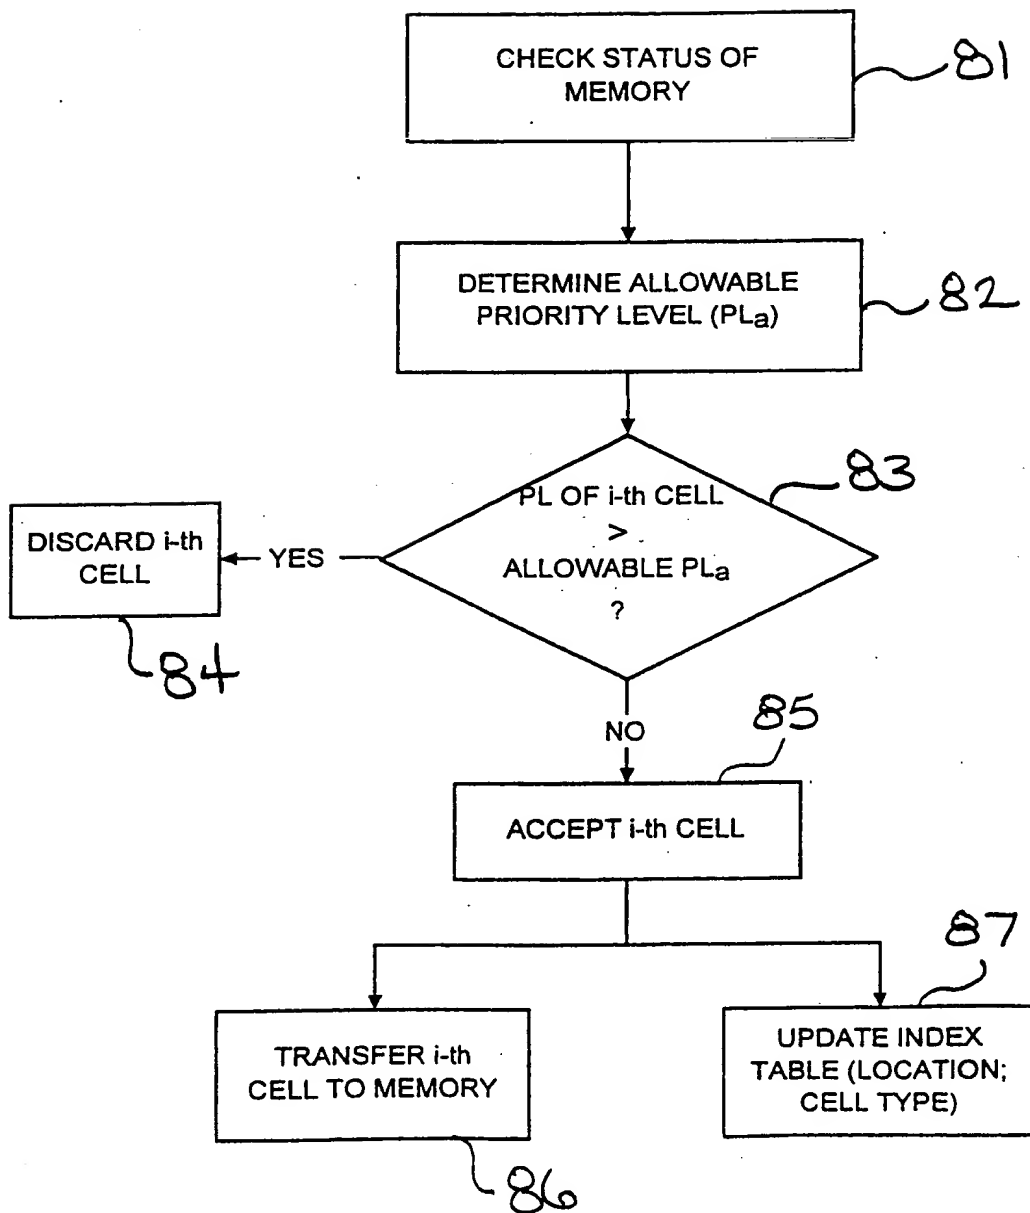


FIG. 3

FIG. 4



5/18




FIG. 5

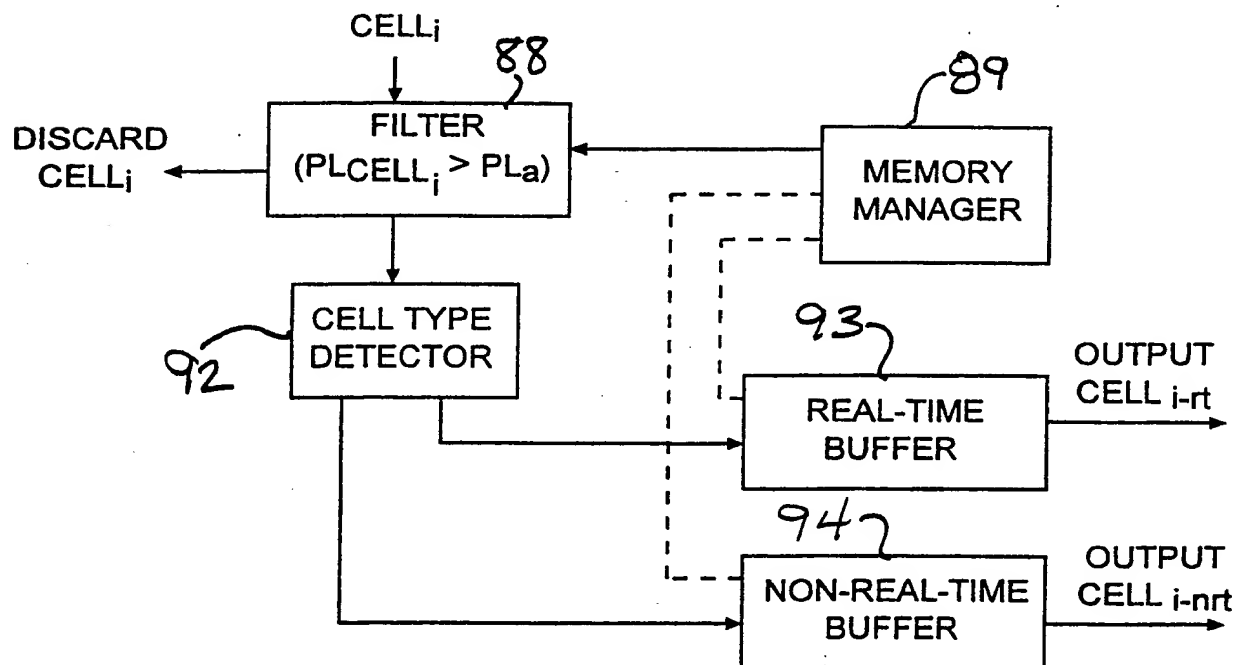


FIG. 6




FIG. 7

6/18




FIG. 8

7/18

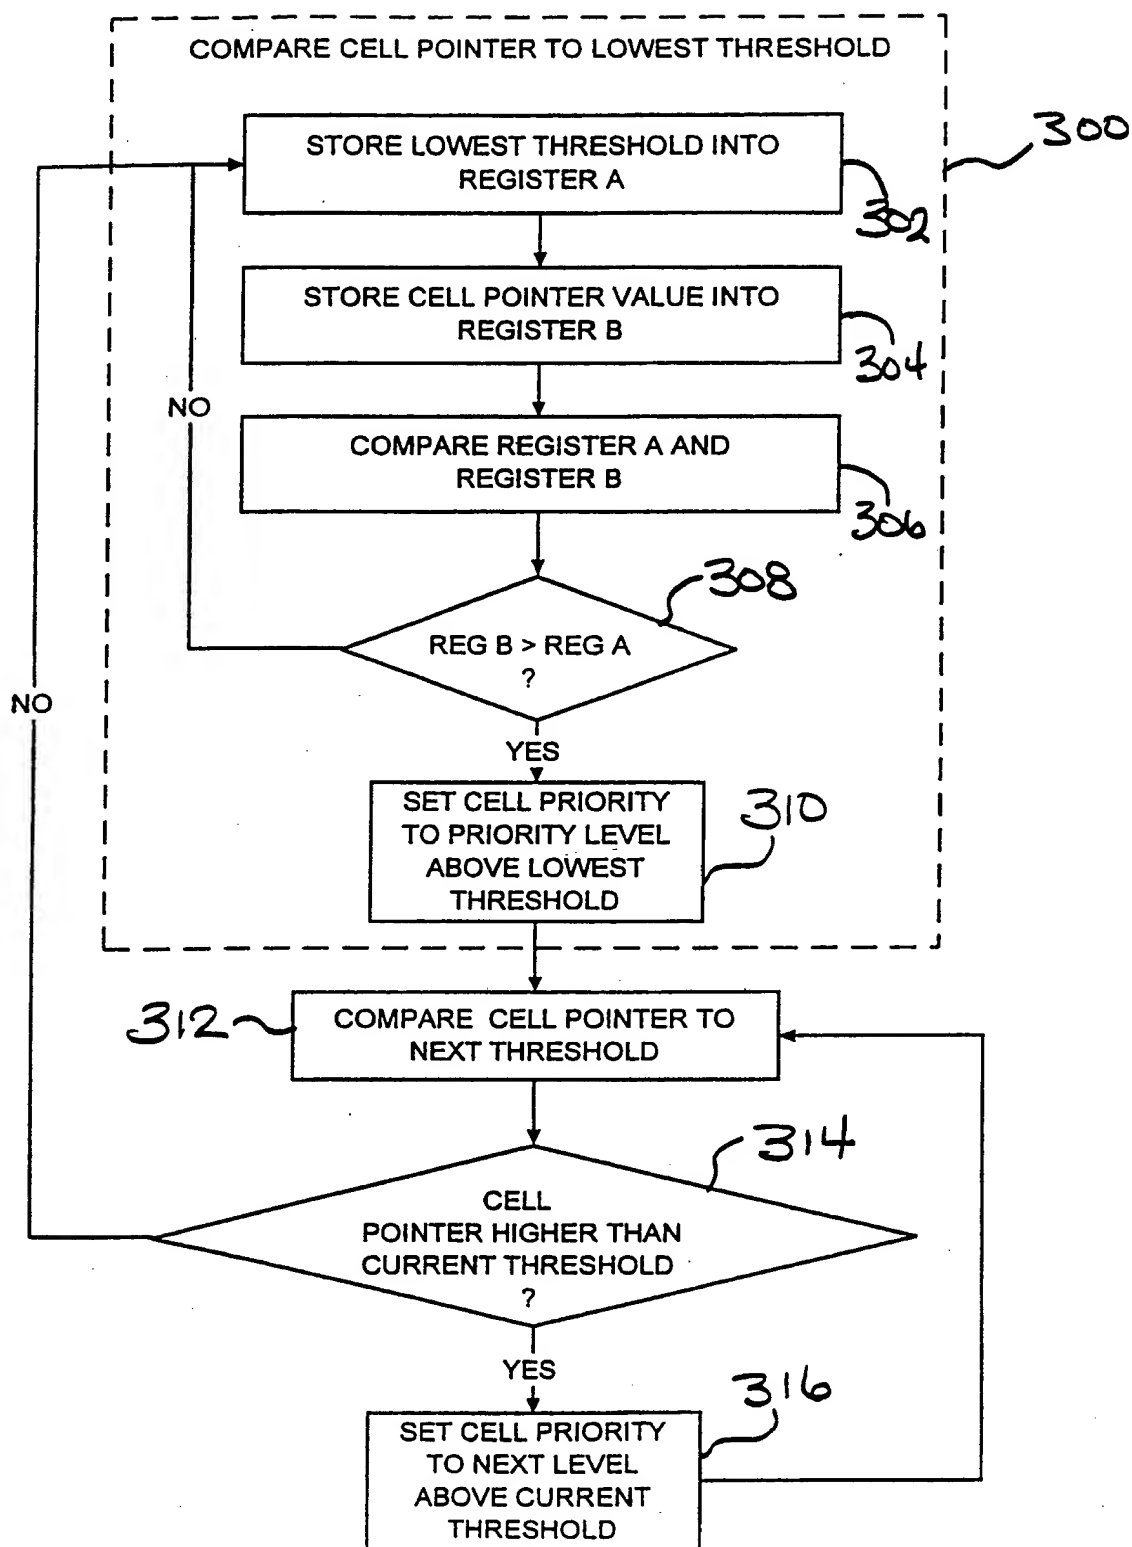


FIG. 9

8/18

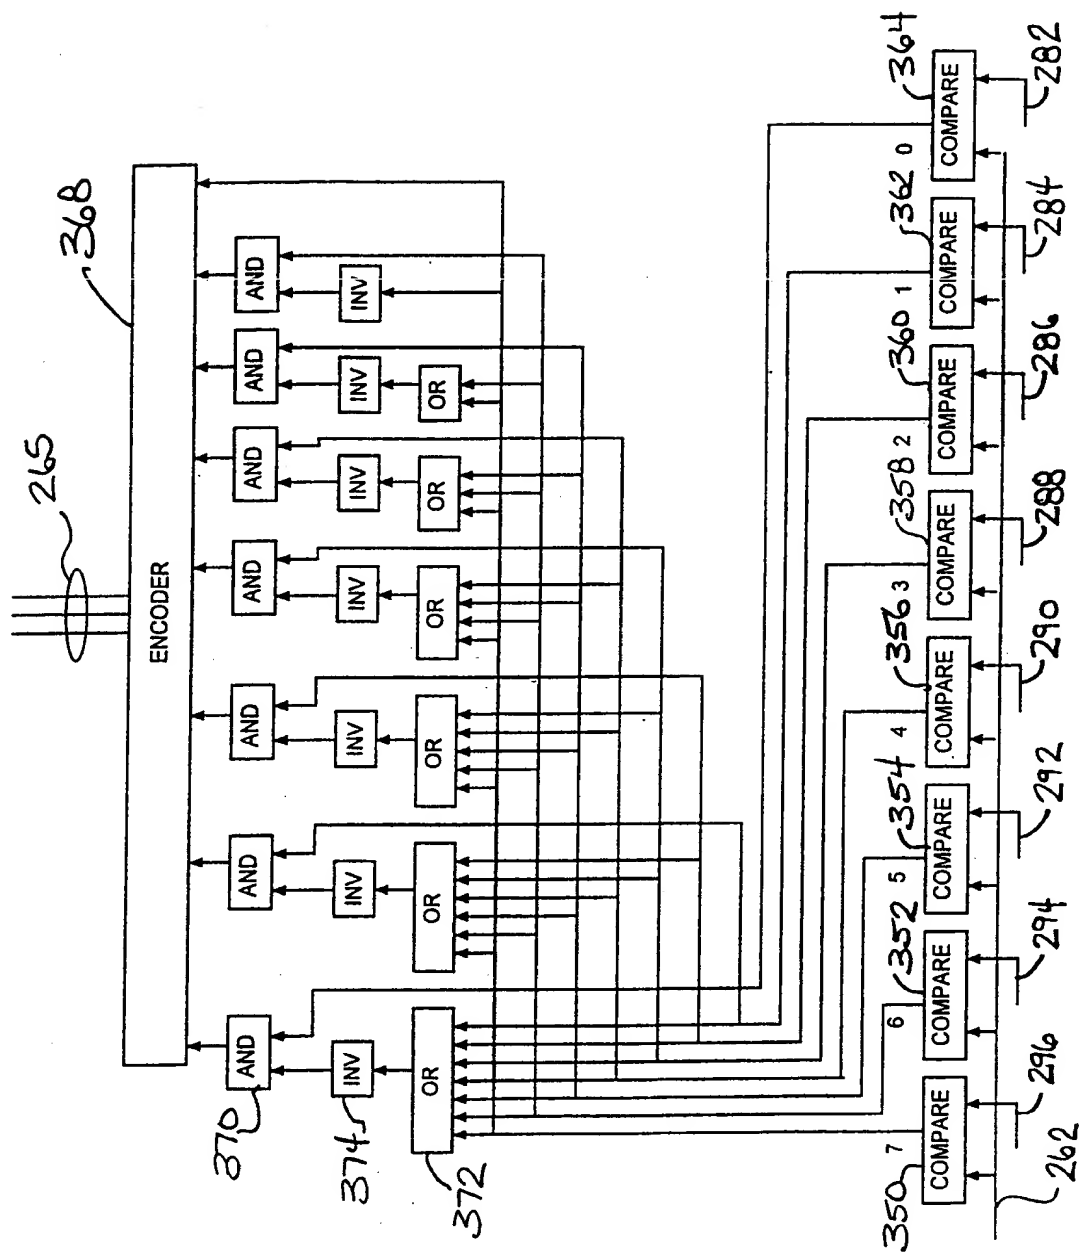


FIG. 10

9/18




FIG. 11

10/18




FIG. 12

11/18

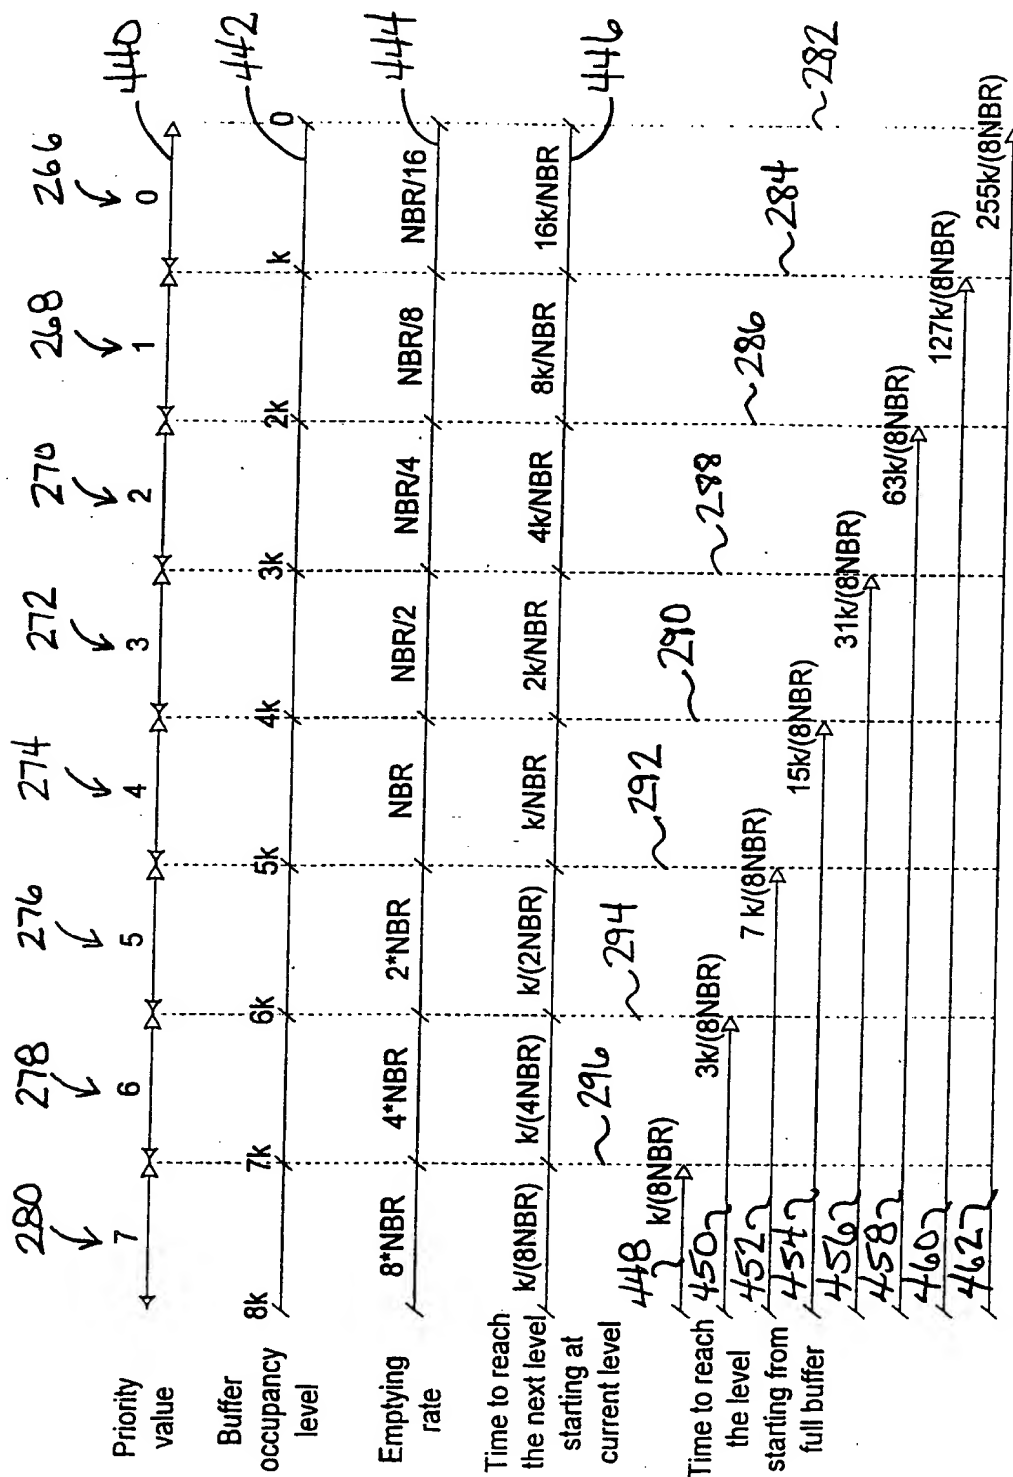


FIG. 13

12/18




FIG. 14

13/18

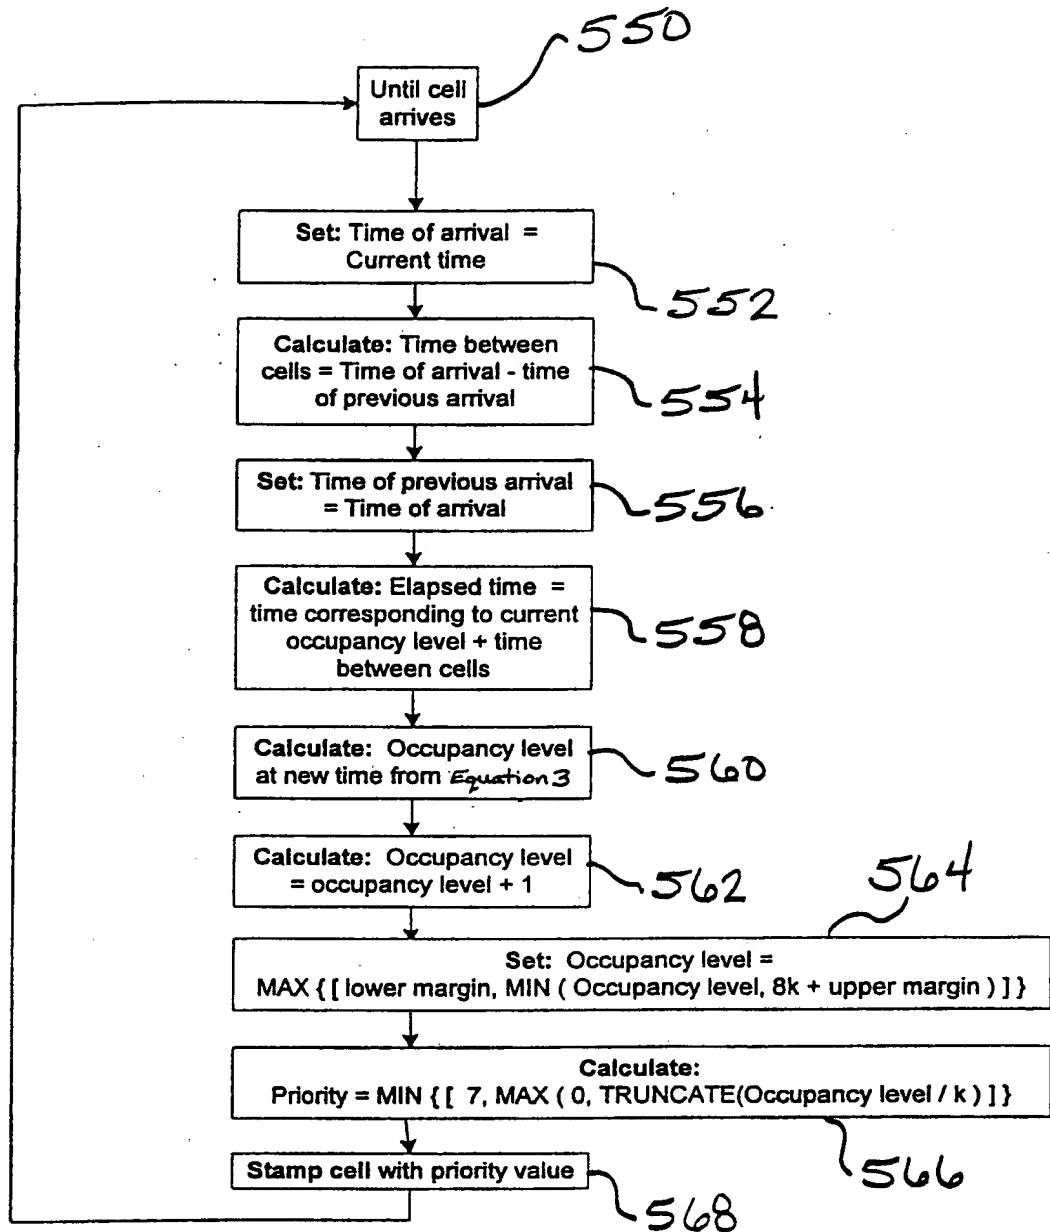


FIG. 15




FIG. 16




FIG. 17

16/18




FIG. 18




FIG. 19

18/18




FIG. 20




FIG. 21

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 99/00232

A. CLASSIFICATION OF SUBJECT MATTER

IPC7: H04Q 11/04, H04L 12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC7: H04Q, H04L, H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5838663 A (ANWAR ELWALID ET AL), 17 November 1998 (17.11.98), column 22, line 47 - column 28, line 36, abstract --	1-38
A	US 5434848 A (PHILIP F. CHIMENTO, JR. ET AL), 18 July 1995 (18.07.95), column 2, line 7 - line 49, claim 1, abstract --	1-38
E,A	US 5909443 A (ALINE FICHO ET AL), 1 June 1999 (01.06.99), column 4, line 45 - line 59; column 5, line 28 - column 8, line 53, claims 1-14, abstract --	1-38

☒ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"I" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

9 December 1999

Date of mailing of the international search report

15-12-1999

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Johanna Lindqvist/cs
Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 99/00232

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0805578 A2 (FORE SYSTEMS, INC.), 5 November 1997 (05.11.97), column 2, line 11 - column 3, line 19; column 6, line 33 - column 13, line 4, claims 1-33, abstract --	1-38
A	EP 0743803 A2 (KABUSHIKI KAISHA TOSHIBA), 20 November 1996 (20.11.96), column 4, line 9 - column 5, line 52, claims 1-36, abstract --	1-38
A	US 5313454 A (LIONEL A. BUSTINI ET AL), 17 May 1994 (17.05.94), column 4, line 35 - column 6, line 37, claims 1-47, abstract -- -----	1-38

INTERNATIONAL SEARCH REPORT
Information on patent family members

02/11/99

International application No.
PCT/FI 99/00232

Patent document cited in search report			Publication date	Patent family member(s)		Publication date
US	5838663	A	17/11/98	CA	2181206 A	25/01/97
				EP	0756403 A	29/01/97
				JP	9046336 A	14/02/97

US	5434848	A	18/07/95	NONE		

US	5909443	A	01/06/99	NONE		

EP	0805578	A2	05/11/97	JP	10065710 A	06/03/98
				US	5935213 A	10/08/99

EP	0743803	A2	20/11/96	JP	8265333 A	11/10/96
				US	5889761 A	30/03/99

US	5313454	A	17/05/94	AU	3797093 A	08/11/93
				EP	0634068 A	18/01/95
				JP	7508378 T	14/09/95
				WO	9320637 A	14/10/93
